TREND ReportPack Guide

TREND 3.6

ReportPacks 4.0

DeskTalk Systems, Inc.

TREND ReportPack Guide
Publication No. RPR361-2000JUN16

DeskTalk Systems, Inc.
19191 S. Vermont Avenue
Ninth Floor
Torrance, CA 90502
(310) 630-1000

Contents

TREND

Figures

Tables

Preface

TREND is collection of related tools, which can provide detailed information about the health and performance of a network. TREND helps a network administrator troubleshoot problem areas, maximize performance, and conduct historical and trending analysis.

The TREND reports that are described in this guide are the key to successful network administration. This guide can help you accurately analyze the data they contain. For information about generating these reports, see the *TREND User's Guide*.

Organization of This Guide

The chapters of this guide are organized as follows:

- Chapter 1 describes installation of ReportPacks, dependencies between ReportPacks, and selection of targets/elements for Snapshot reports.
- Chapter 2 describes the uniform layout of the reports and lists the eight reports common to most ReportPacks.

- Chapters 3 analyzes selected reports from each ReportPack to give you guidelines for interpreting the information provided in any one report type.
- Chapters 4–12 describe each ReportPack in detail, including descriptions of each report, each chart within a report, data sources, supported devices, and metric and GOS calculations.
    - Chapter 4 (Frame Relay Report Pack) focuses on the Frame Relay switch and PVCs.
    - Chapter 5 (LAN and WAN ReportPacks) focuses on LAN and WAN devices.
    - Chapter 6 (Router ReportPack) focuses on the various router interfaces.
    - Chapter 7 (Bay Router ReportPack) focuses on the Bay router interfaces.
    - Chapter 8 (Cisco Router ReportPack) focuses on the Cisco router interfaces.
    - Chapter 9 (RMON ReportPack) focuses on the various segments of Ethernet networks, Token Ring networks, or both.
    - Chapter 10 (RMON2 ReportPack) focuses on data collected from RMON2 probes.
    - Chapter 11 (TREND Database ReportPack) focuses on the historical and predicted usage levels of the TREND database and transaction log.
    - Chapter 12 (System ReportPack) focuses on statistics about CPU utilization, paging and swapping, file system usage, virtual and physical memory, cache hits and misses, and grade of service specifically for the Empire SystemEDGE agent.
- Appendix A describes the property file formats for various ReportPacks.

Related Documents

This guide is designed to help you understand each of the ReportPacks, and use them to accurately analyze the data they contain. Use this guide in conjunction with the following guides:

- *TREND User's Guide*
- *TREND Installation Guide*
- *Guide to Viewing TREND Reports*
- *Guide to Building TREND Reports*

Online User Manual Availability

TREND is shipped with online-viewable user manuals in Portable Document Format (PDF), along with the Adobe® Acrobat® Reader. These documents and the Adobe Acrobat Reader are available prior to installation in the Documents directory of the installation CD. You need to install the Adobe Acrobat Reader to view these documents online.

TREND document and associated file names are as follows:

Document Name	File Name
TREND Installation Guide	TREND*xx*_InstallationGuide.pdf
TREND User's Guide	TREND*xx*Users_Guide.pdf
TREND ReportPack Guide	TREND*xx*ReportPack_Guide.pdf
TRENDweb Installation and Administration Guide	TRENDweb*xx*_InstallGuide.pdf
Guide To Building TREND Reports	TRENDweb*xx*_BuildersGuide.pdf
Guide To Viewing TREND Reports	TRENDweb*xx*_ViewersGuide.pdf

where *xx* is the product version number. (For example, 32 refers to version 3.2).

These documents are also installed onto your hard drive in the docs directory in DPIPE_HOME during installation.

You can download the latest TREND manual updates from the DeskTalk Systems Customer FTP site. Look at the document number (on the manual's title page) to determine if you have the latest version of manual:

IG	*TREND Installation Guide*
UG	*TREND User's Guide*
RPR	*TREND ReportPack Guide*
TW-IAG	*TRENDweb Installation and Administration Guide*
BTR	*Guide To Building TREND Reports*
VTR	*Guide To Viewing TREND Reports*

Product version (for example, 32 means Version 3.2.

Publication date in *yyyymmmdd* format.

TW-IAG 32-2000MAR06

Call DeskTalk Systems Customer Support for download instructions.

TREND

TREND

TREND

1 ReportPack Administration

This chapter discusses the following:

- ReportPack dependencies
- ReportPack installation
- Using the property import feature
- Selecting deferred targets or elements for a snapshot report

ReportPack Dependencies

Some ReportPacks use reports from other ReportPacks. Therefore, ensure that you install the required ReportPack(s), if any, for the ReportPack you are using. Table 1-1 lists the dependencies for each ReportPack.

Table 1-1: ReportPack Dependencies

ReportPack	Dependencies
Frame Relay	None
LAN/WAN	None
Router	LAN/WAN
Bay Router	Router, LAN/WAN
Cisco Router	Router, LAN/WAN
RMON	None
RMON2	None
TREND Database	None
System	None

Installing ReportPacks

The V4.0 Bundled ReportPack suite comprises the following ReportPacks:

- Frame Relay
- LAN_WAN Connectivity
- Router

- Bay Router
- Cisco Router
- RMON
- RMON2
- TREND_Database (replaces the Database ReportPack)
- System

When you installed TREND V3.6, the installation program placed the ReportPacks sources in the `{DPIPE_HOME}/packages` directory. For more information about installing the ReportPacks from those sources, see "Running Autopilot" on page 1-10.

Note: In this chapter, {DPIPE_HOME} represents the directory into which TREND has been installed. By default, this is the filesystem with the most space on UNIX and C:\TREND on Windows NT.

If you have downloaded the ReportPacks from the DeskTalk web site, then you have to uncompress the ReportPack archive file. For more information about unzipping the compressed archive file, see "Unzipping the ReportPack Sources" on page 1-3.

Unzipping the ReportPack Sources

Note: Before you uncompress the ReportPack sources, ensure that you remove or rename any previous versions of the ReportPacks that reside in the {DPIPE_HOME}/packages directory.

To uncompress the ReportPack archive file:

1. If you do not have an unzip utility installed, you can download one by entering the following URL in your web browser:

http://www.cdrom.com/pub/infozip/UnZip.html

2. Unzip the ReportPack archive file to the `{DPIPE_HOME}/packages` directory on the system where you have TREND installed, and ensure that the option that preserves the directory and path names is selected.

3. Verify that you now have a folder for the ReportPack under the `{DPIPE}\packages` folder.

4. Use AutoPilot to install the ReportPacks. For more information about using AutoPilot, see “Running Autopilot” on page 1-10.

ReportPack Directory Structure

The V4.0 Bundled ReportPack suite has a directory structure, which did not exist for previous versions. A ReportPack may have different installation options. For example, the ReportPack directory may contain a folder for the ReportPack, its demo version, and the ReportPack’s upgrade. Figure 1-1 shows an example of what the directory structure for the System ReportPack might look like.

```
TREND {DPIPE_HOME}
  |
  |------- packages
  |           |
  |           |------- System
  |           |           |
  |           |           |------- System.ap
  |           |           |------- System_Demo.ap
  |           |           |------- UPGRADE_System.ap
```

Figure 1-1: System ReportPack—Example Directory Structure

Types of ReportPack Installations

You can perform the following types of ReportPack installations: new, upgrade, or demo.

New Installation

A new installation installs a new version of the ReportPack. For more information about installing a new ReportPack, see "Running Autopilot" on page 1-10.

Note: **Before you install a ReportPack, ensure that no previous versions of it are installed. If there is a previous version, remove it using Autopilot, or you can install an upgrade version of the ReportPack if one is available.**

Upgrade Installation

An upgrade installation upgrades an existing version of a ReportPack. You can install an upgrade version of the ReportPack (for example, `UPGRADE_System <Version 4.0>`) if you have a V3.5 version of the ReportPack. For more information about installing an upgrade, see "Running Autopilot" on page 1-10.

The upgrade installation checks for a preceding version of the ReportPack before it completes the upgrade. The upgrade installation preserves any existing data; however, it installs new RPTs, scripts, and TEEL files, and it performs any required database schema updates. Autopilot will show that the V4.0 ReportPack is installed and that the upgrade package was installed.

Note: The upgrade installation is only available for pre-V3.5 RMON and RMON2 ReportPacks. If you have other pre-V3.5 ReportPacks, remove them, and install the V4.0 ReportPacks available with TREND V3.6. When you upgrade from these pre-V3.5 ReportPacks to V4.0 ReportPacks, all of your existing data is no longer supported.

Note: Do not remove the upgrade package unless you also intend to remove the ReportPack.

Note: The old TRENDsheet (.qss) and TRENDgraph (.qgr) reports are no longer supported. You can manually remove the directories containing these files from {DPIPE_HOME}/reports directory. This ensures that they will not be displayed in the list of available TRENDview PRO reports.

Upgrading from a V3.4.x RMON or RMON2 ReportPack

To upgrade from a V3.4.x RMON or RMON2 ReportPack:

1. Edit the `install.pkg` file in the following directory:

   `{DPIPE_HOME}/packages/<ReportPack>/UPGRADE_<ReportPack>.ap`

2. Comment out the dependency and version lines, for example:

```
### Identification
report_pack:UPGRADE_RMON
version:4.0

### Check that an earlier version is installed
#dependency:RMON
#version:3.5
```

3. Using AutoPilot, add the UPGRADE_RMON ReportPack. For more information about running AutoPilot, see “Running Autopilot” on page 1-10.

Note: Before you can add the RMON2 ReportPack, you must run an upgrade procedure for the V3.6 binary set to operate correctly. For more information about running the upgrade procedure, see “Running the RMON2 ReportPack Upgrade Procedure” on page 1-7.

Running the RMON2 ReportPack Upgrade Procedure

For the RMON2 ReportPack only, you must run an upgrade procedure for the V3.6 binary set to operate correctly.

The following bug is being fixed by this upgrade:

- Incorrect translation of DECNET and AIX addresses in host and matrix RMON2 tables

The `rmon2_341_upgrade.exe` program performs following changes:

- Adds a new column `address_type` to the `dsi_protocol` table.
- Populates the `address_type` column in the `dsi_protocol` table.
- Creates and populates the `dsi_version_info` table.
- Updates the `dsi_table_key`, `nl_srcaddress`, and `nl_dstaddress` columns in RMON2 key tables.

Note: The `rmon2_341_upgrade.exe` **program does not change the** `table_keyxxx` **column in data tables; the changes are made only to the column in Key Tables.**

Some Key Tables (Ghost Key Tables) are affected by the upgrade.

Note: For UNIX systems, DeskTalk recommends that you back up the Ghost Key Tables before running the `rmon2_341_upgrade.exe` **program. For more information about backing up the Ghost Key Tables and restoring key tables, see Chapter 5 of the *TREND Installation Guide*.**

The following Key Tables (Ghost Key Tables) are affected by upgrade:

- `Gmon2_nlhostdata_` —applicable to users who collected the `RMON2 NL HOST` table with the 3.41 version of the collector.
- `Gmon2_alhostdata_` —applicable to users who collected the `RMON2 AL HOST` table with the 3.41 version of the collector.
- `Gmon2_nlmatrixdata_` —applicable to users who collected the `RMON2 NL MATRIX` table with the 3.41 version of the collector.
- `Gmon2_almatrixdata_` —applicable to users who collected the `RMON2 AL HOST` table with the 3.41 version of the collector.

Note: Ensure that all the rmon2 collections are turned off while the upgrade procedure is running.

To run the RMON2 upgrade procedure:

1. Run the following program:

```
%DPIPE_HOME%\bin\rmon2_341_upgrade.exe
```

2. Using AutoPilot, add the UPGRADE_RMON2 ReportPack. For more information about running AutoPilot, see "Running Autopilot" on page 1-10.

Troubleshooting the RMON2 Upgrade Procedure

If an error occurs during the upgrade procedure, examine the`%DPIPE_HOME%\tmp\trend.log` file for details.

During Key Table Upgrade stage, the upgrade program tries to recreate indexes on the key tables. This process requires a large amount of `temp_db` and database log. If the upgrade fails during index recreation because there is not enough space, you must recreate the index manually after additional space is allocated.

To recreate index, run the following command:

```
isql -Udsi_dpipe -Pdsi_dpipe
1> create unique index K2 on ghost_key_tablename
2> (dsi_control_table_keyid, dsi_table_key)
3> go
1> quit
```

Demo Installation

A demo installation installs a ReportPack that has example reports. Each ReportPack provides a demo package (for example, `System_Demo.ap`) which you can install. When you install the demo ReportPack, example report files are installed in the following directory:

```
{DPIPE_HOME}/reports/Demo/<ReportPack>subdirectory
```

In this directory, the report files are visible to all TRENDweb servers and all TRENDview PRO clients.

Note: You can view the demo reports (.sdt files) without a TREND system using TRENDview Pro. If you prefer to do this, do not use Autopilot to install the demo, use TRENDview PRO to open the .sdt files. For more information about TRENDview PRO, see the *Guide to Viewing TREND Reports*.

Running Autopilot

The installation of a ReportPack using Autopilot:

- Creates all the data tables, key tables and stored procedures.
- Adds lines to the `trend_timer.sched` file, which schedules the ReportPack's `trend_proc` files to run.
- Installs the TRENDweb reports.

To run Autopilot:

1. Ensure that you are logged in as the `trendadm` user.

2. Click the **Start** button, and then point to **Programs**. Point to the folder that contains TREND, and then click **TREND**. The **TREND** window is displayed.

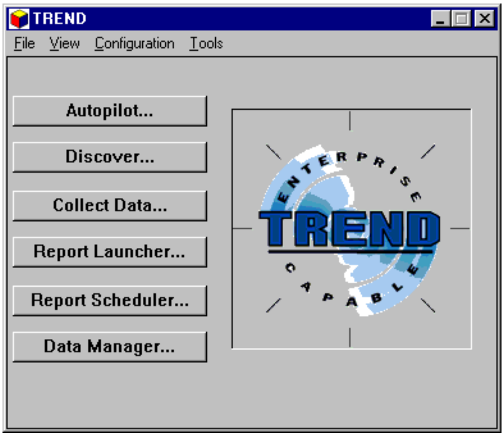

3. Click **Autopilot**. The **Autopilot** dialog box is displayed.

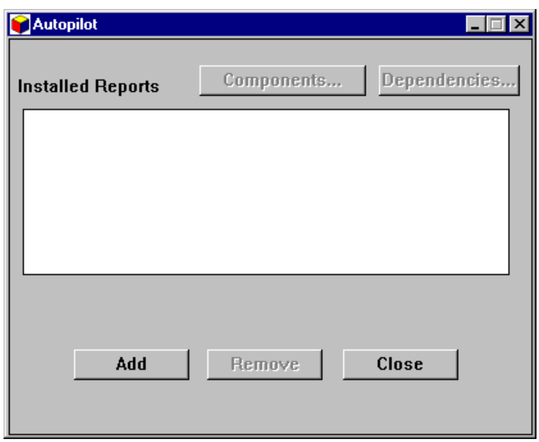

4. Click **Add**. The **TREND Report Packs** dialog box is displayed.

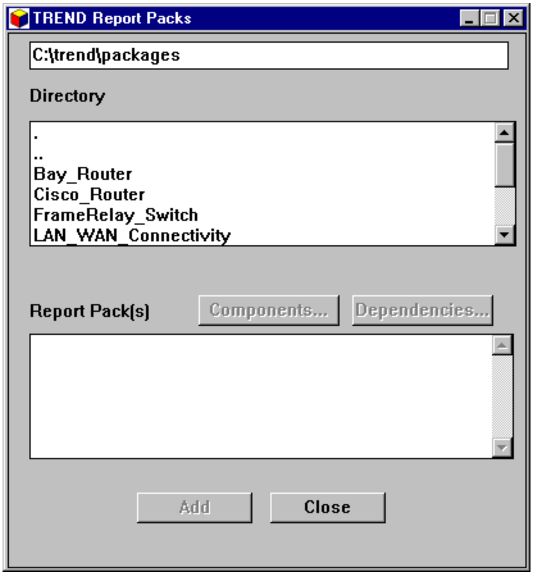

5. Double-click the name of the ReportPack (for example, **System**) in the **Directory** list. The **TREND ReportPacks** dialog box displays the available ReportPack(s).

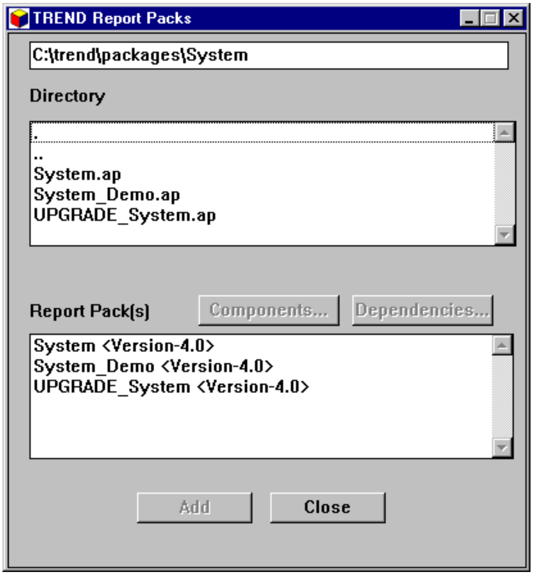

6. Click the name of the ReportPack you want to install (for example, **System** <**Version 4.0**>) in the **ReportPack(s)** list.

   You can choose one of the following types of installations:

   - **New**—Installs a new version of the ReportPack (for example, `System <Version 4.0>`)
   - **Demo**—Installs a demonstration package with example report files (for example, `System_Demo <Version 4.0>`).
   - **Upgrade**—Installs an upgraded version of an existing ReportPack (for example, `UPGRADE_System <Version 4.0>`).

   For more information about the ReportPack options, see “Types of ReportPack Installations” on page 1-5.

7. Select the ReportPack option you want, and click **Add**.

TREND begins the installation. Ignore any “`No data found to load for table xxxx`” messages; these are normal.

TREND will display a message when the installation is complete to ask if you want to begin to collect data.

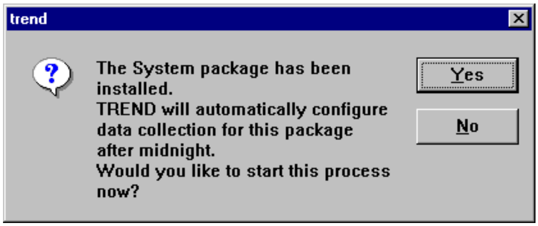

8. Click **No** unless you have previously run IP Discovery. For more information about IP Discovery, see Chapter 5 of the *TREND User’s Guide*.

Using the Property Import Feature

The property import feature allows you to import user-provided property information that relates to the objects that a ReportPack reports on. Specifically, the property import feature allows you to set up the Executive Summary reports so that they only display data for devices belonging to a specific customer.

This feature is only available for the following V4.0 Bundled ReportPacks:

- LAN/WAN Connectivity
- Router
- Cisco Router
- Frame Relay
- Bay Router

This feature will be available for other selected ReportPacks at a later date.

Note: **You can use the property import feature after you install the ReportPack. However, use of the property import feature is not mandatory. If you do not need to use it, you can ignore the subsequent sections.**

This feature is useful for service-provider organizations who want to provide segmented reporting to different users or customers. For example, suppose Thunderbolt Inc. has offices in Boston, Pittsburgh, and in Los Angeles. The TRENDweb Administrator can configure the key tables after the ReportPack installation so that users in the Boston, Pittsburgh, and Los Angeles offices will be treated as separate customers, seeing only the devices belonging to their office. Thus, when a user displays an Executive Summary report, the report only shows performance statistics for the devices relevant to the user's office. The user's office is displayed at the top of the Executive Summary report. However, when the TREND Administrator logged in as `trendadm`, he would see all the offices shown under the Customer Name category and a special All category, which would show the same statistics for all of the devices in the network. Figure 1-2 shows an example of an Executive Summary report that displays devices specifically for the Thunderbolt Inc. customers.

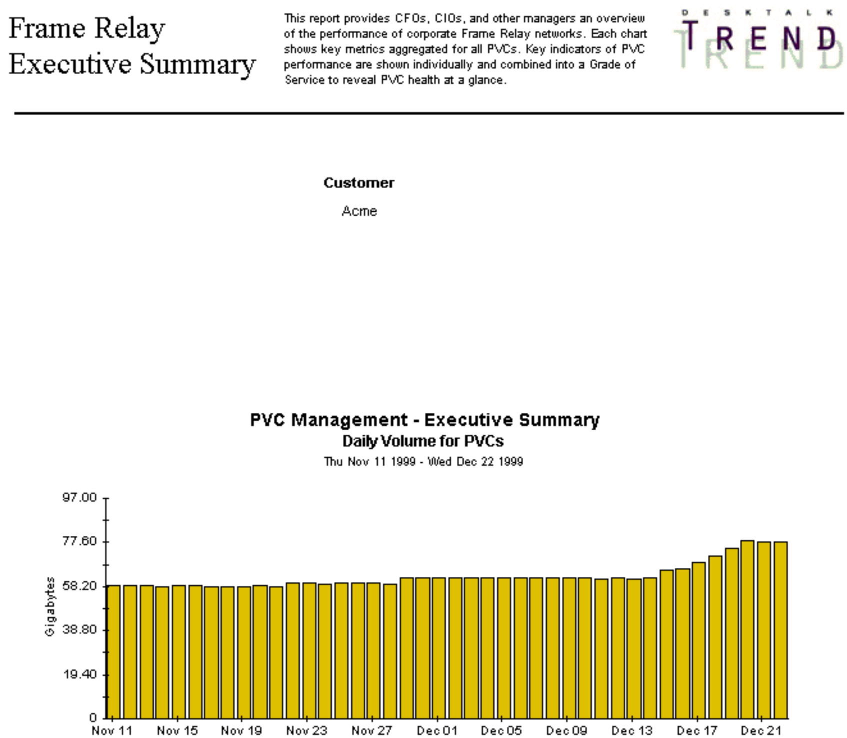

Figure 1-2: Devices Displayed for a Specific Customer

The property import feature allows you to import `CustomerID` and `CustomerName` assignments for the objects being reported on.

The `CustomerID` assignments are used by:

- The ReportPack when it processes data to provide aggregated summary information. This information is used in Executive Summary reports.
- TRENDWeb security functions to determine which data to provide to which users.

How Property Import Works

Figure 1-3 shows a high-level view of the architecture of the property import feature.

Figure 1-3: Architecture of the Property Import Feature

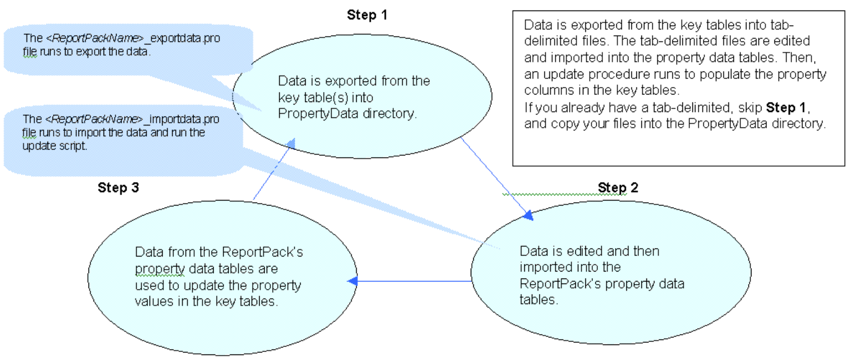

Data is exported from the key tables to a property file, which is a tab-delimited ASCII file. You can generate it using the `<ReportPack>_exportdata.pro` process, or you can create it using an application like Microsoft Excel. Then, you edit the data and import the file into the ReportPack's property data tables. The data is used to update property information contained in the key tables. The format of the the property file is uniquely defined for each ReportPack. For more information about the property file formats for the bundled ReportPack, see "Property File Formats" on page A-1.

You must place the property file into the designated `SourceDirectory`; the ReportPack will process the file from this directory. TREND uses a configuration file called a TEEL file to determine how to find property import files. The ReportPack uses the TEEL file to locate a new property file or files nightly and imports the contents of those found. It overwrites existing property assignments and insert new property assignments.

The TEEL file for each ReportPack is installed into the `{DPIPE_HOME}/lib` folder by Autopilot during the ReportPack's installation. The TEEL file has the following file name format:

`<ReportPack>_Property.teel`

For example, the TEEL file for the Frame Relay ReportPack would be `Frame_Relay_Property.teel`.

Note: Although the examples in the TEEL file show that the SourceDirectory resides in the {DPIPE_HOME}/packages directory, DeskTalk recommends that you create and use a new directory instead. This will avoid any possible problems when you install future ReportPack upgrades in the {DPIPE_HOME}\packages directory.

If you change the SourceDirectory location, as DeskTalk recommends, you have to edit the ReportPack's TEEL file to define the `SourceDirectory` and `SourceDisposition`.

Defining the SourceDirectory

The `SourceDirectory` directive tells TREND where to look for the property import files for a ReportPack. It specifies the full path name to the directory and a file name specification that uses wildcard characters. The wildcard characters allow TREND to match on file names.

The syntax of the `SourceDirectory` directive differs depending on the operating system you are using. The TEEL file provides examples of the syntax for both Windows NT and UNIX.

To specify the `SourceDirectory` directive:

1. Edit the ReportPack's TEEL file using a text editor like Notepad.
2. Ensure that the `SourceDirectory` directive has the correct path and file name specification to match the incoming property import files.

Defining the SourceDisposition

The `SourceDisposition` directive tells TREND what to do with the imported property file after processing it. DeskTalk recommends that you use the `move` disposition to direct TREND to move the file to an archive directory.

To specify the `SourceDisposition` directive:

1. Edit the ReportPack's TEEL file using a text editor like Notepad.
2. Ensure that the directory you specify for the `SourceDisposition` exists with write access.

Sample Property File

Figure 1-4 shows a sample property file (`Frame_Relay_Property.dat`) for the Frame Relay ReportPack.

```
Sanjose 1.20 4620000 512000  2 Acme   2 Unassigned RegionSupport
atlanta 1.24 4620000 512000  2 Acme  -2 Unassigned RegionMarketing
atlanta 1.42 4620000 512000  2 Acme  -2 Unassigned RegionMarketing
lincoln 1.19 4620000 512000  2 Acme  -2 Unassigned RegionFinance
lincoln 1.79 4620000 512000  2 Acme  -2 Unassigned RegionSupport Server
orlando 1.30 4620000 512000  2 Acme  -2 Unassigned RegionFinance
orlando 1.50 4620000 10240002 Acme  -2 Unassigned RegionFinance Server
sanjose 1.40 4620000 512000  1 ABC   -2 Unassigned RegionMarketing
portland1.21 4620000 512000  1 ABC   -2 Unassigned RegionAccounting
torrance1.89 4620000 512000  1 ABC   -2 Unassigned RegionSupport
torrance1.90 4620000 512000  1 ABC   -2 Unassigned RegionMarketing
portland1.30 4620000 512000  1 ABC   -2 Unassigned RegionSales
torrance1.88 4620000 512000  1 ABC   -2 Unassigned RegionSales
   ❶       ❷     ❸       ❹    ❺  ❻    ❼        ❽          ❾
```

Figure 1-4: Sample Property File

1. `dsi_target_name`
2. `dsi_table_key`
3. `port_speed`

4. `pvc_cir`
5. `cust_id`
6. `customer_name`
7. `region_id`
8. `region_name`
9. `dsi_descr`

The property file has been populated for all columns, except `region_id` and `region_name`, and is also set up to display only devices belonging to Acme or ABC customers. For a description of each property column, see Table A-2.

Displaying Data for Specific Customer Devices in Reports

To set up an Executive Summary report to show devices by customer:

1. Generate a property file. For more information about generating or creating a property file, see "Exporting Data" on page 1-21.
2. After you generate or create the property file, edit it to define the `cust_id` and `customer_name` values in the property table. (By default, these are assigned a value of -2 and are consider unassigned.)

   For example, the property file for the Frame Relay ReportPack might look like the following after you generate it:

   ```
   Sanjose  1.20       -2  Unassigned Customer-2 Unassigned Region
   atlanta  1.24       -2  Unassigned Customer-2 Unassigned Region
   atlanta  1.42       -2  Unassigned Customer-2 Unassigned Region
   lincoln  1.19       -2  Unassigned Customer-2 Unassigned Region
   lincoln  1.79       -2  Unassigned Customer-2 Unassigned Region
   orlando  1.30       -2  Unassigned Customer-2 Unassigned Region
   orlando  1.50       -2  Unassigned Customer-2 Unassigned Region
   ```
3. Supply a unique integer value for `cust_id`. However, you cannot use -2, -1, or 0.

For example, suppose the `SanJose` and `atlanta` targets belonged to Thunderbolt Inc. and the `lincoln` and `orlando` targets belonged to Belle Express. After you edit the property file, it would look like the following:

```
Sanjose  1.20          1   Thunderbolt Inc.-2  Unassigned Region
atlanta  1.24          1   Thunderbolt Inc.-2  Unassigned Region
atlanta  1.42          1   Thunderbolt Inc.-2  Unassigned Region
lincoln  1.19          2   Belle Express   -2  Unassigned Region
lincoln  1.79          2   Belle Express   -2  Unassigned Region
orlando  1.30          2   Belle Express   -2  Unassigned Region
orlando  1.50          2   Belle Express   -2  Unassigned Region
```

4. Ensure that the User ID value you assign using the TRENDweb User Administration application matches the `cust_id` value that you assigned in the property datapipe for each customer. For more information about how to create user accounts, see the *TRENDweb Installation and Administration Guide*.

Note: If the cust_id does not match the value of User ID, then you may not see your reports.

5. Import the property file using the instructions in "Importing Property Data" on page 1-24.

Note: If a new PropertyImport changes existing CustomerID assignments, TREND will begin using the new assignments immediately for all subsequent aggregations. However, aggregations that were previously processed will not be reprocessed based on the new assignments.

Exporting Data

To export data, you must generate a property file containing the data you want the importing process to use by doing either of the following:

- Use the property datapipe to generate the property file(s).
- Create your own property files.

Generating Property Files

To generate a property file:

1. From a command prompt, change directory to the `{DPIPE_HOME}/scripts` directory.
2. Locate the ReportPack's procedure file and run it using the following command:

   `trend_proc -f <ReportPackName>_exportdata.pro`

   This process calls `trend_export`, which reads a definition file to determine which columns to export from the database. It exports a template property file. The file contains a row for each object that is currently discovered by the ReportPack. If no objects are discovered, then the file contains no rows. Therefore, this utility is only useful after discovery and at least one polling cycle has been completed. All exported columns come from the key table(s). `Trend_export` generates a tab-delimited property file and place it in the following folder:

   ```
   DPIPE_HOME/packages/<ReportPackName>/<ReportPackName>.ap
   /PropertyData
   ```

After you generate the file, you can edit it manually or using an application like Microsoft Excel and then export the data to a tabbed-delimited file.

Creating Your Own Property File

It does not matter what application you use to create the property file. However, the property file must be an ASCII, tab delimited file, and the format of the file must comply with one of the column specifications in "Property File Formats" on page A-1.

An example property file is provided for each ReportPack. It is located in the following directory:

```
{DPIPE_HOME}/packages/<ReportPack>/<ReportPack>.ap/
PropertyData
```

If you create your own tab-delimited file, you must:

- Use the file formats described in “Property File Formats” on page A-1.
- Use all columns shown in “Property File Formats” on page A-1 in the same order.
- Do **not** use quotes in your file.
- Do **not** use NULL values in your file.
- Ensure that the property data columns of `dsi_target_name` and `dsi_table_key` match the values of `dsi_target_name` and `dsi_table_key` in the key table.

  For example, if you have a `dsi_target_name` of `router1` with a `dsi_table_key` value of 1.23 in your property file, there must be a matching `dsi_target_name` of `router1` with a `dsi_table_key` value of 1.23 in the key table. Otherwise, a new row and keyid value will be inserted into the key table.
- Set the values of the `*_id` values to -2, which is the default value, if you are not using some or all of them.

  For example, if you are not using `region_id`, set the value for `region_id` in your file to be -2.
- Use the following file naming convention:

  ```
  <ReportPackName>_Property.dat
  ```
- Place the property data files in the following folder:

  ```
  {DPIPE_HOME}/packages/<ReportPackName>/<ReportPackName>.ap
  /PropertyData/Archive
  ```

To create your own property file:

1. Create a tab-delimited file using an application like Microsoft Excel.
2. Edit or add properties.
3. Export the properties as a tab-delimited file.

Importing Property Data

Property import is handled automatically by the system if the files are present in the designated `SourceDirectory`.

To import property data before the scheduled nightly process:

1. From a command prompt, change directory to the {DPIPE_HOME}/scripts directory.
2. Locate the TREND procedure file, and run the procedure file using the following command:

   `trend_proc -f <ReportPackName>_importdata.pro`

This will truncate the existing property data table and import the new property data into the property data table using `ee_collect`.

The update script will update the property columns in the key table only if an entry exists in the property data table that has the same keyid value in both data table and key table.

Removing the ReportPack

Note: **You cannot remove a ReportPack if other ReportPacks have a dependency on that ReportPack.**

To remove the ReportPack:

1. Click the **Start** button, and then point to **Programs**. Point to the folder that contains TREND, and then click **TREND**. The **TREND** window is displayed.

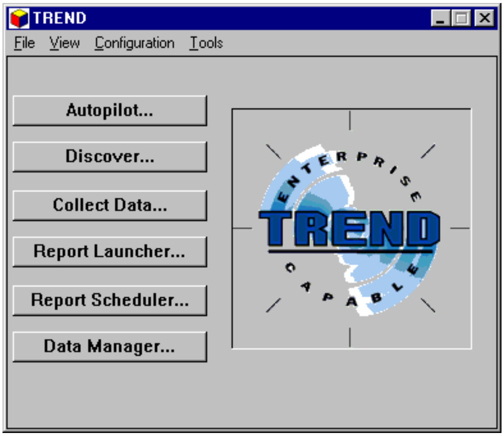

2. Click **Autopilot**. The **Autopilot** dialog box is displayed.

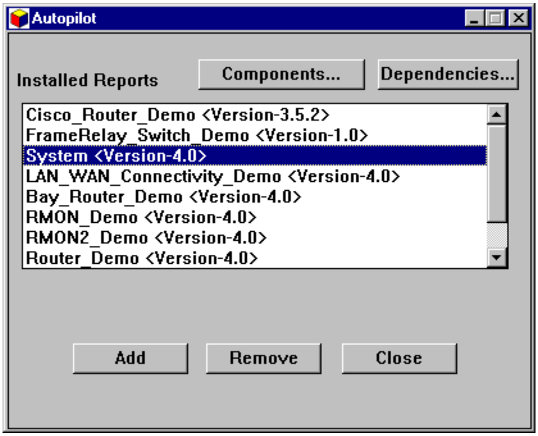

3. Select the ReportPack (for example, `System <Version-4.0>`), and click **Remove**. The **Remove ReportPack** dialog box is displayed.

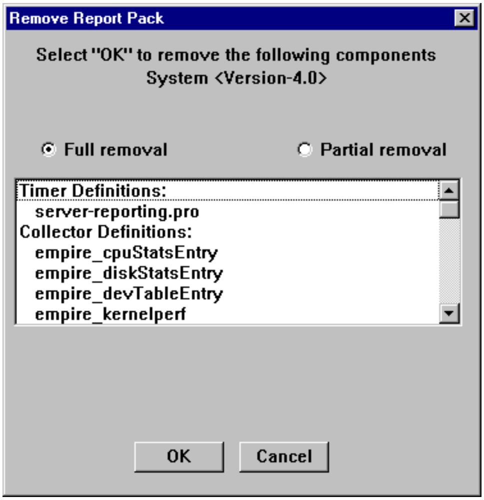

4. Click **Full Removal**, and click **OK**.

This deletes all the data tables and modifies `trendtimer.sched` as appropriate.

Selecting Deferred Targets or Elements for a Snapshot Report

A snapshot report is one of the eight standard reports in a ReportPack. It allows you to view statistics for a selected target or element.

To select a target or element for a Snapshot report:

1. Select the Snapshot report from the Report Listing area of the TRENDview Pro window.

   The **Choose one or more Targets or Elements** dialog box is displayed.

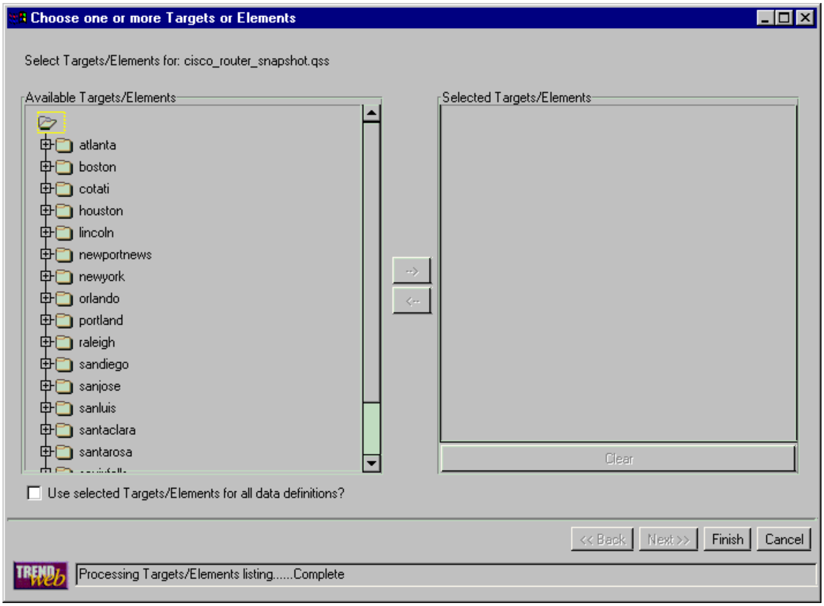

2. Select the target or element you want to view from the **Available Targets/ Elements** box, and click the arrow to move it to the **Selected Targets/ Elements** box.
3. Select any other elements you want to view, or if you are finished, click **Next**.

The Snapshot report is displayed with performance data for the selected target or element.

2 Overview

This chapter:

- Describes the uniform layout of the reports in all TREND ReportPacks (see "Report Layout" on page 2-2).
- Lists the eight reports that are common to most ReportPacks and briefly describes the content and purpose of each report (see "Standard Reports" on page 2-4).
- Identifies the metrics (volume, throughput, and so on) that are reported for each enterprise technology area (Frame Relay Circuits, Routers, and so on) and gives an overview of how each metric is calculated (see "Metric Classes" on page 2-6).
- Explains the concepts of forecasting, baselining, ranking, and frequency distribution as they are applied in TREND reports (see "Forecasting" on page 2-8, "Baselines" on page 2-10, "Ranking" on page 2-13, and "Frequency Distribution" on page 2-14).

Report Layout

All TREND reports have a uniform layout that makes the report easy to understand. Figure 2-1 shows a sample report.

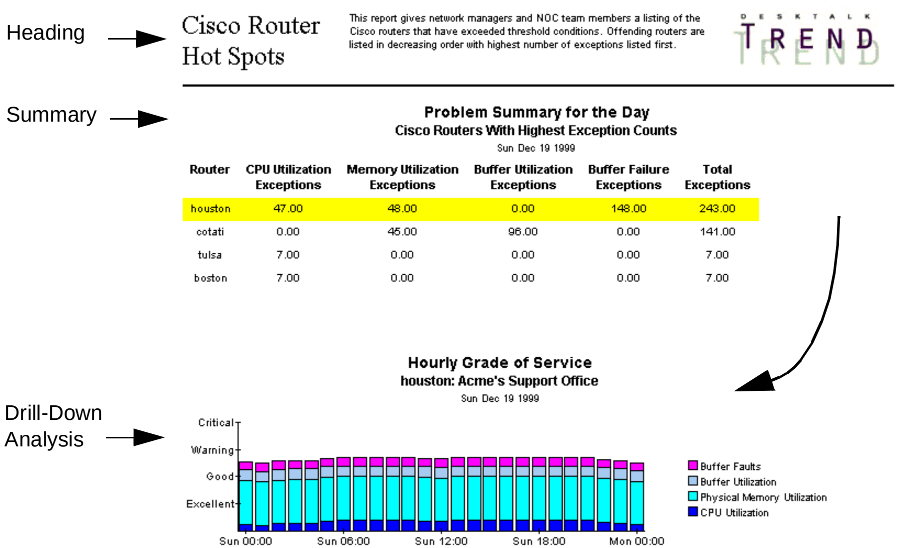

Router	CPU Utilization Exceptions	Memory Utilization Exceptions	Buffer Utilization Exceptions	Buffer Failure Exceptions	Total Exceptions
houston	47.00	48.00	0.00	148.00	243.00
cotati	0.00	45.00	96.00	0.00	141.00
tulsa	7.00	0.00	0.00	0.00	7.00
boston	7.00	0.00	0.00	0.00	7.00

Figure 2-1: TREND Report Layout

The report components are as follows, although not all reports contain drill-down analysis:

- **Heading**. The heading identifies the report and gives a high-level description of the report contents. Figure 2-1 shows the Cisco Router Hot Spots report.
- **Summary**. Depending on the report type, the summary information is presented as a table, as shown in Figure 2-1, or a graph. It highlights the macro issues associated with the report type. For the Cisco Router Hot Spots report shown in Figure 2-1, the Problem Summary for the Day table identifies the Cisco routers, listed in descending order by the greatest number of total exceptions, that are exhibiting the most problematic performance.

  In this example, problematic performance is reported as exception conditions in the areas of CPU, memory, buffer utilization, and buffer faults. An exception condition occurs when the predefined threshold for a metric (for example, 95% CPU utilization) is exceeded. You can modify the default thresholds supplied in the TREND ReportPacks to suit your environment.
- **Drill-Down Analysis**. The drill-down analysis components are presented as a combination of tables, graphs (line and area), and charts (bar, stacked bar, and pie). They provide a detailed analysis of the individual metrics for the selected managed element over the appropriate time period (preceding day or 6-week baseline period). In Figure 2-1, tabbed drill-down components enable you to examine CPU utilization, memory utilization, buffer utilization, and buffer faults for the selected router over each hour of the preceding day. These multiple, related charts bring together the information you need for decision and action. Note that not all reports have drill-down analysis.

  If you view the report with the TRENDview PRO Client, you can select a row in the summary chart by drilling-down (double-clicking) on it, which highlights the selected row in yellow. The content of the drill-down components changes to reflect the metrics for the selected summary row (in this case, the Houston router). For instructions on using the TRENDview PRO Client, see the *Guide To Viewing TREND Reports*.

Standard Reports

A report consists of a summary chart and one or more detailed drill-down charts, graphs, or tables that support the summary data. Table 2-1 lists the eight standard reports each TREND ReportPack provides

Table 2-1: Standard Reports

Report	Description
Executive Summary	Provides high-level graphs that aggregate key metrics for all managed elements within a technology (for example, all frame relay PVCs or all routers). Executives can review the daily volume summary and, if desired, look into more detailed information such as grade of service, hourly volume compared to baseline, and total exception count.
Hot spots	Identifies specific areas of concern where threshold values have been exceeded (for example, excessive utilization or error counts). Managed elements in this report are ranked by total number of exceptions. Each exception is linked to multiple drill-down components that bring together the information needed to investigate potential network trouble spots, pinpoint the source of the problem, and determine possible resolutions.
Forecast	Delivers a view into the future so that proactive corrections can be made to prevent service problems. Provides a ranked list of all elements forecast to be within 90 days of reaching a threshold exception condition. Reveals the number of days-to-threshold (DTT) and adds baseline utilization values and 90-day projections for accurate forecast of budget requirements and capital equipment expenditure justification. Each element is linked to multiple drill-down components which show trends over the baseline period.

(1 of 2)

Table 2-1: Standard Reports

Report	Description
Capacity Planning	Details the most overutilized and underutilized managed elements in the same report, which indicates opportunities for load balancing to improve service levels without additional investment. Enables the network manager to optimize existing resources and deliver improved service without increasing the budget. Can be used to expose expensive, underutilized resources that are targets for budget savings. Drill-down components examine element performance over time.
Top Ten	Highlights the top volume contributors and the elements with the most severe problems. Includes all key metrics for each of the top ten managed elements. Highlights volatile elements, as significant behavior changes are often indicators of service-level problems in the making. Lists current rank, previous rank, and change in rank, which enables the network manager to assess the severity of network changes.
QuickView	Displays the top ten problem elements (ranked by GOS score) and drill-down charts appropriate to the technology (for example, utilization, errors, and volume compared to baseline).
Snapshot	Includes the same information as the QuickView report, but the element is selected from a pick list that is displayed when the report is invoked. Enables the network manager to view the status of any element immediately, not just the top ten problem elements.
Service Level Management	Shows response time and availability plotted against contracted service levels. Shows whether or not service levels have been achieved for the previous day, month, or other time period.

(2 of 2)

For a selective description of the content of these reports for the various ReportPack technologies, see "Analyzing TREND Reports" on page 3-1. Subsequent chapters provide detailed descriptions of the standard reports for each of the ReportPacks.

Metric Classes

Table 2-2 lists report metrics. One or more of the metrics can be included in each report.

Table 2-2: Metric Classes

Metric	Description
Volume	Amount of information that flows through a managed element. For example, for Frame Relay, volume is the total inoctets + outoctets, in Gigabytes, for all Permanent Virtual Circuits (PVCs).
Utilization	The amount of information actually flowing through a managed element divided by the capacity of the device. Utilization is given in terms of bandwidth for Frame Relay PVCs and in terms of percentage of CPU time, memory, and buffer space for LAN or WAN interfaces.
Throughput	Amount of information that passes through a managed element completely and correctly over time. The metric is given in millions of bits per second (Mbps).

(1 of 3)

Table 2-2: Metric Classes

Metric	Description
Congestion	Measures the number of frames that experience congested areas on the network in both forward and backward directions. Indicates the level of impediments to throughput.
Grade of Service (GOS)	Weighted scores for selected metrics that, together, add up to a number in one of the following ranges that signifies a GOS score. The following is an example of GOS grading based on a scale of four points: **Range** — **GOS Score** 0.0 - 1.0 — Excellent 1.1 - 2.0 — Good 2.1 - 3.0 — Warning 3.1 - 4.0 — Critical
Latency	Amount of time it takes for information to pass through a managed element (or the network).
Availability	Time a managed element is functional divided by the entire time period (uptime for the managed element).

(2 of 3)

Table 2-2: Metric Classes

Metric	**Description**
Frequency Distribution	Obtained by dividing a sample’s frequency in a defined class (bucket) by the total number of samples. For example, frequency distribution can be depicted with a pie chart, showing the time that an element spends in a specified utilization range over a day.

(3 of 3)

Each metric must be interpreted in terms of the technology area to which it applies. For example, utilization is shown in terms of bandwidth for Frame Relay PVCs and in terms of CPU, memory, and buffer utilization for routers.

For an explanation of how the metric is applied to a specific technology, see the subsequent chapter that describes that technology’s ReportPack.

Forecasting

Forecasts are used in the Capacity Planning and Forecast reports, where resource usage is forecast for 30-, 60-, and 90-day periods and the number of days until a specified usage threshold (DTT or days-to-threshold forecasting) is reached.

The 30-, 60-, and 90-day forecasts answer the following question: What will the metric value (for example, percent utilization) be for a resource in 30, 60, or 90 days?

The days-to-threshold forecast answers the following question: In how many days will the resource being forecast be at the threshold value?

A default threshold value is supplied with each TREND ReportPack, which you can change to suit your installation.

TRENDsum forecasting calculations are based on a least squares linear regression, which means that the calculations attempt to draw a straight line through a set of data points collected over time, as shown in Figure 2-2.

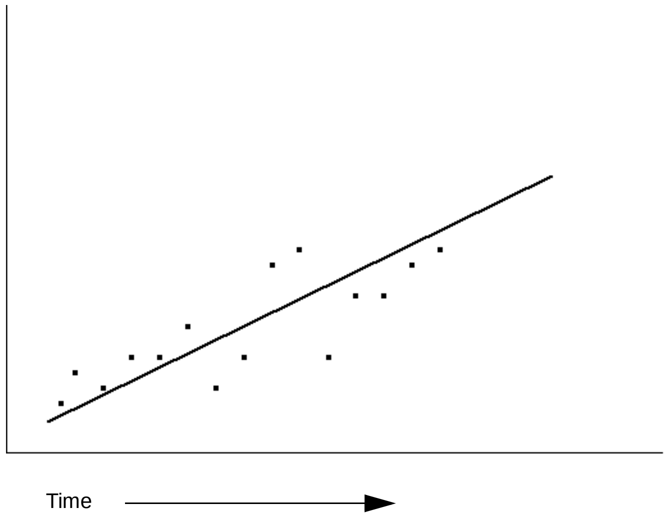

Figure 2-2: Forecasting—Least Squares Linear Regression

Each data point represents a daily 95th percentile value for the element property being forecast. The 95th percentile is the value below which 95% of the samples taken during the day fall. You want to look at peak utilization for planning purposes, so a daily 95th percentile value, which is a near-maximum value, is used rather than a simple average. The 95th percentile better characterizes the shape of the data by taking into account the measurement variations (spikes) that occur during the day; averaging, on the other hand, tends to smooth out the effect of measurement variation, thereby giving a false impression of capacity.

For a DTT calculation, you have the following cases:

- You will hit the threshold in the calculated number of days.
- You are already beyond the threshold (a negative value).
- You are never going to hit the threshold (the line slopes downward, away from the threshold value).
- You are a long distance from the threshold. If the threshold is too far away, the forecast is meaningless.

If the slope is negative (decreasing, sloping downward), the result returned for DTT is a null value (displayed as Not Applicable in reports). Likewise, if the threshold is more than 1,000 days away, the result returned is null and displayed as Not Applicable.

For formulas that TRENDsum uses to compute all forecasts, see the section on the TRENDsum utility in Appendix A of the *TREND User's Guide*.

Baselines

A *baseline* provides a basis for comparing recent performance with past performance.

TREND ReportPacks use a *rolling baseline* period of 42 days (6 weeks) to compute baselines. In a rolling baseline, the baseline table contains *n* rows; each row represents the same statistic (or statistics) computed for the immediately preceding number of days in the baseline period. For example, if you use a 42-day baseline period, the rolling daily baseline is computed as follows:

1. At the end of day 1, the baseline statistic is computed from all the sample data collected for day 1 and stored in a summary row for day 1. For example, assume you are computing the average value of an element property to use as a baseline. If 96 samples are collected for day 1 (one sample every 15 minutes),

the values for all samples are totaled and divided by 96 to compute the average for the day 1:

Day	Baseline Statistic
1	a_1

2. At the end of day 2, the baseline statistic is recomputed from all the sample data collected for days 1 and 2 and stored in the summary row for day 1, thus *replacing the existing day 1 value*. For example, if 96 samples are collected for day 2, those 96 values and the 96 values collected on the previous day (day 1) are totaled and divided by the sample count (now 192—96 for day 1 plus 96 for day 2) to arrive at value a_{2}.

Day	Baseline Statistic
1	a_2

3. In the same fashion, the baseline statistic is recomputed at the end of day 3 from all the sample data collected for days 1, 2, and 3 and stored in the summary row for day 1, again replacing the existing day 1 value. At this point, the value (a_3) for the baseline summary row represents a 3-day aggregate:

Day	Baseline Statistic
1	a_3

4. The same processing is repeated at the end of each day until a baseline value has been recalculated for each day of the rolling baseline period (42 days, in

this example). Thus, at the end of 42 days, a single baseline value has been computed based on data collected on the previous 42 days:

Day	Baseline Statistic
1	a_{42}

This day 1 baseline value, because 42 days of data is used to compute it, is significantly more refined than a value where only one day's samplings are summarized to compute the baseline statistic for day 1.

5. When the number of days in the rolling baseline period is reached, processing continues in the fashion described above. However, the rolling baseline value that is computed from the preceding 42 days of samples is appended to the table. The existing baseline is no longer replaced. Thus, rows are added to the rolling baseline summary table as follows:

Day	Baseline Statistic	Computed from
1	a_1	Days 1 thru 42
2	a_2	Days 2 thru 43
3	a_3	Days 3 thru 44
4	a_4	Days 4 thru 45
n	a_n	Days *n* thru (*n*+41)

The Executive Summary reports in the TREND ReportPacks compare a resource usage for each hour of the day to an hourly baseline that has been computed for the resource. The concept for computing an hourly baseline is the same as the one for computing a daily baseline. However, the hourly baseline table has one row for each hour of each day computed for a sliding 42-day baseline period as described above.

Ranking

Ranking lets you chart the volatility of a given metric for a given element. For example, suppose you want to use utilization and congestion among 100 PVCs to answer the following questions:

- Which ten PVCs rank the highest in terms of utilization and congestion on any given day?
- How often does the same element appear in the top ten in terms of utilization and congestion?
- Are there significant variations in element rank from day to day (delta rank)?

The ranking table has one row per day for each managed element that is being ranked and the following six columns:

- Timestamp (for day1, day2, day3, and so on)
- ID of the metric that is being ranked
- Daily summary value of the metric being ranked
- Rank of that metric (1, 2, 3 and so on)
- Rank for the same metric on the previous day
- Delta rank—This item represents the absolute value of the difference in the metric's value between yesterday and today. A large delta rank indicates a volatile element, which you should probably investigate. A relatively small delta rank indicates stable element performance.
- Previous delta rank

Ranked metrics appear in the Top Ten reports in the TREND ReportPacks.

Frequency Distribution

The frequency distribution of a sample is obtained by dividing the sample's frequency in a defined class (bucket) by the total number of samples.

For example, the In Utilization metric for the Frame Relay ReportPack is calculated for each sample taken during the day for the selected PVC. After the In Utilization percentage is recorded, a value of 1 is assigned to one of the defined buckets (0-20%, 21-40%, 41-60%, 61-80%, 81-100%, and >100%) depending on the recorded percentage. If the In Utilization percentage for one sample is 42%, then a value of 1 is assigned to the 41-60% bucket.

The values for the buckets are totaled for all samples taken during the day, then the total is divided by the number of samples taken, which results in a bucket-total/sample-count ratio. That ratio is then multiplied by 100 to give a percentage of samples for the day whose In Utilization value falls within that bucket. Table 2-3 shows a sample calculation of frequency distribution for the In Utilization metric.

Table 2-3: Sample Calculation for Frequency Distribution

Input Utilization Buckets	**Frequency**	**Frequency Distribution**
0-20%	7	7/30 = 23.3%
21-40%	12	12/30 = 40%
41-60%	7	7/30 = 23.3%
61-80%	3	3/30 = 10%
81-100%	1	1/30 = 3.7%
>100%	0	0/30 = 0%
Total	**30**	**100%**

Figure 2-3 shows a sample of how TREND depicts frequency distribution.

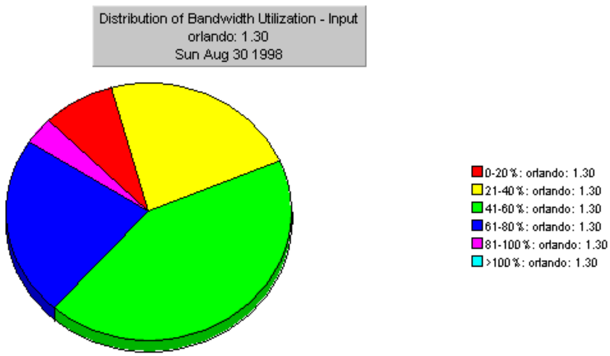

Figure 2-3: Sample Chart for Frequency Distribution

TREND

3 Analyzing TREND Reports

ReportPacks typically contain the following reports:

- Executive Summary
- Hot Spots
- Forecast
- Capacity Planning
- Top Ten
- QuickView
- Snapshot
- Near Real Time
- Service Level Management

Each report generally shows the same type of information for all ReportPacks. For example, the Forecast report for the Router ReportPack forecasts router resources for the next 30, 60, and 90 days and identifies the days to threshold (DTT) for selected router resources. The Forecast report for the Frame Relay ReportPack provides the same information for selected Frame Relay PVCs.

The way metrics are computed differ from one ReportPack to the next. For example, utilization for the Frame Relay ReportPack is given in terms of PVC bandwidth. Conversely, utilization for the Cisco Routers ReportPack is given in terms of CPU, memory, and buffer consumption.

This chapter analyzes selected reports from each ReportPack. For example, the Executive Summary report from the LAN ReportPack and the Forecast report from the Frame Relay ReportPack are analyzed in detail. The intent of this approach is to give you guidelines for interpreting the information provided in any one report type. If you apply the detailed explanation of a report contained in this chapter along with the ReportPack-specific metrics in subsequent chapters, you should be able to interpret the reports in any ReportPack.

Executive Summary Report

The Executive Summary report provides high-level graphs aggregating key metrics for all managed elements within a technology (for example, all LAN interfaces, all Frame Relay PVCs, all WAN Interfaces, or all Routers). The report is structured for the corporate executive, who can review the daily volume summary and, if desired, look into more detailed information such as Grade of Service (GOS), hourly volume compared to baseline, and total exception count. The intent of the report is to provide executives with insight into the usage and return on investment of their network resources.

For example, the following components comprise the LAN Connectivity Executive Summary:

- Daily Volume chart (see “Daily Volume Chart” on page 3-3).
- Hourly Volume chart (see “Hourly Volume and Hourly Baseline Charts” on page 3-4).
- Hourly Baseline chart (see “Hourly Volume and Hourly Baseline Charts” on page 3-4).
- Hourly Grade of Service chart (see “Hourly Grade of Service Chart” on page 3-5).

- Hourly Bandwidth Utilization chart (see “Hourly Bandwidth Utilization Graph” on page 3-6).

The following sections describe these charts in more detail. For complete information about the LAN Connectivity ReportPack, see “LAN and WAN ReportPacks” on page 5-1.

Daily Volume Chart

Figure 3-1 shows the Daily Volume bar chart, which compares yesterday’s total volume of traffic to the daily volumes observed over the baseline period (which is the 42 days from November 9 through December 19 in this example). In this example, *yesterday* is December 19.

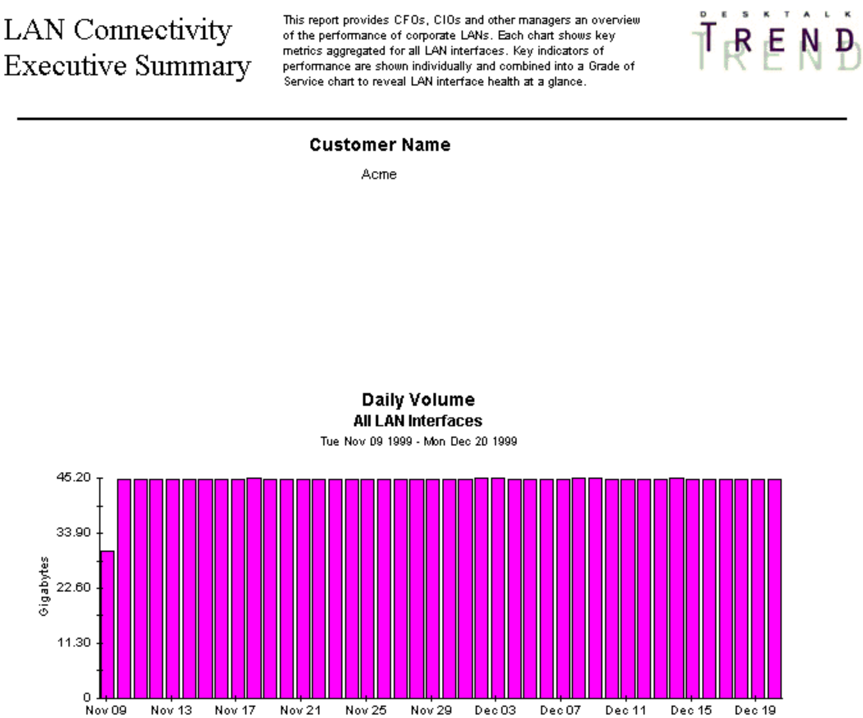

Figure 3-1: Executive Summary— Daily Volume Chart (LAN)

It also puts the most recent volume information in context with the long-term historical trend that is shown over the baseline period.

Note: The TRENDweb Administrator can configure a ReportPack installation, so that the Executive Summary reports will only display data for devices belonging to a specific customer. In Figure 3-1**, the report displays information for only the devices belonging to the Acme company. For more information about this feature, see** “Displaying Data for Specific Customer Devices in Reports” on page 1-20**.**

Hourly Volume and Hourly Baseline Charts

Figure 3-2 shows the Hourly Volume bar chart, which details the total volume of traffic for each hour (yesterday) and focuses the executive’s attention on the growth and changes over the last 24 hours.

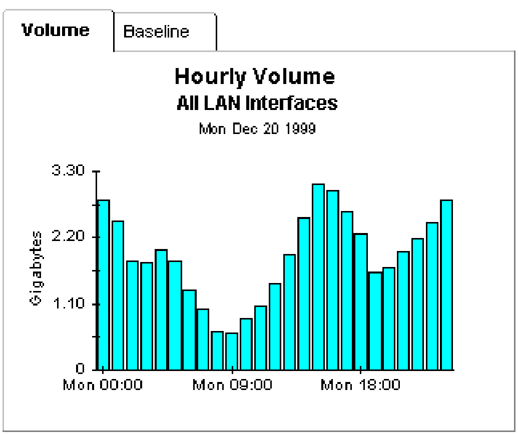

Figure 3-2: Executive Summary— Hourly Volume Chart (LAN)

Click on the **Baseline tab** to see the Hourly Baseline bar chart. Figure 3-3 shows the Hourly Baseline chart.

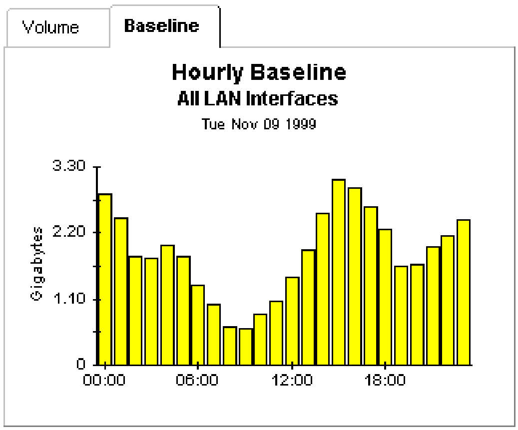

Figure 3-3: Executive Summary— Hourly Baseline Chart (LAN)

This chart shows the average volume by hour of day over the baseline period. The hourly baseline is a much more valuable indicator than a simple average or a daily average baseline. Since traffic typically varies by hour of day—for example, the traffic at 10:00 a.m. is typically much higher than at midnight—TREND calculates individual baseline values for each hour of the day over the baseline period to deliver the most accurate and useful comparisons. Comparing this chart to the Hourly Volume chart will allow the executive to see if yesterday's traffic deviated significantly from the historical average.

Hourly Grade of Service Chart

Figure 3-4 shows the Hourly Grade of Service stacked bar chart, which correlates a number of important service quality metrics appropriate to each technology.

Health scores for each metric are assigned and weighted according to their relative importance and added with other metric scores to calculate a final grade. In the LAN

ReportPack, the key metrics contributing to overall health are Utilization, Discards, and Errors. By assigning a grade and corresponding description (for example, Excellent, Good, Warning, Critical) to each hour of the day, TREND provides an aggregate view of the stability or changes in overall health for the class.

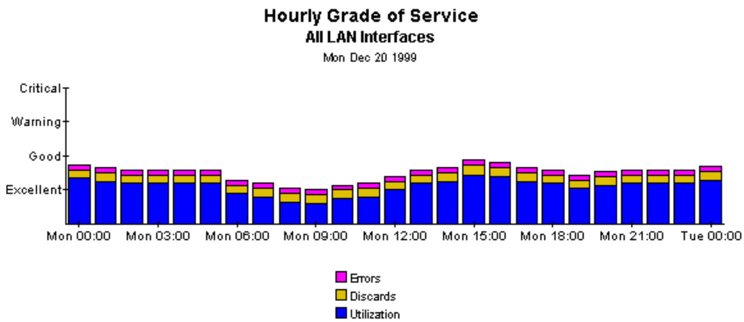

Figure 3-4: Executive Summary—Hourly Grade of Service Chart (LAN)

Hourly Bandwidth Utilization Graph

Figure 3-5 shows the Hourly Bandwidth Utilization line graph, which shows executives volume expressed as a percentage of available bandwidth.

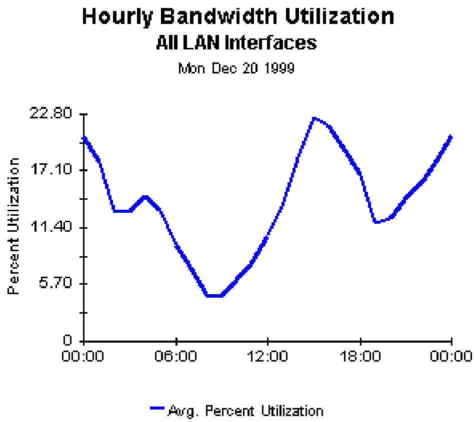

Figure 3-5: Executive Summary—Bandwidth Utilization Graph (LAN)

If the chart shows that current traffic is approaching the top end of currently accessible resources, additional investment may be required to satisfy end user demands.

Forecast Report

The Forecast Report delivers a view into the future so that proactive corrections can be made to prevent service problems. The Forecast report is also a valuable resource to plan network expansion and justify future capital expenditures.

For example, the following components comprise the Forecast report for Frame Relay:

- Estimated Days to Threshold (DTT) summary table (see "Estimated Days to Threshold (DTT) Summary" on page 3-7) and the following drill-down components.
    - Grade of Service chart (see "Grade of Service Drill-Down" on page 3-10).
    - Utilization chart (see "Utilization Drill-Down" on page 3-11).
    - Congestion chart (see "Congestion Drill-Down" on page 3-12).

The following sections describe these charts in more detail. For detailed information about the Frame Relay ReportPack, see "Frame Relay ReportPack" on page 4-1.

Estimated Days to Threshold (DTT) Summary

Figure 3-6 shows the Estimated Days to Threshold (DTT) summary table, which provides a ranked list of all elements that have exceeded or are forecasted to be within 90 days of reaching a threshold exception condition. The DTT is shown along

with the average of the daily P95 utilization over the baseline period and the 90-day forecasted P95 utilization for both in and out utilization.

Estimated Days to Threshold (DTT)
Projected To Exceed Thresholds Within 90 Days
All Values Based On 95th Percentile

Device	Interface/DLCI	DTT In Utilization	Current P95 In Utilization	Projected In Util 90 Days	DTT Out Utilization	Current P95 Out Utilization	Projected Out Util 90 Days
Sanjose	1.20	-3.00	87.96	122.82	-22.00	97.19	131.37
portland	1.30	-220.00	176.03	211.02	-217.00	180.04	211.97
torrance	1.88	-3.00	88.34	122.18	-22.00	94.61	130.31
sanjose	1.40	-3.00	88.49	123.42	-23.00	95.64	128.84
torrance	1.89	-225.00	180.22	209.42	-210.00	175.58	213.80
torrance	1.90	-4.00	87.47	125.44	-22.00	97.31	130.19
portland	1.21	-210.00	178.13	213.26	-218.00	177.84	212.01
orlando	1.50	-1000.00	366.85	377.01	-1000.00	369.21	384.86
orlando	1.30	61.00	57.27	93.89	63.00	58.00	92.79

Figure 3-6: Forecast—Estimated Days to Threshold Table (Frame Relay)

The managed elements are sorted with those representing the most immediate risk listed first. TREND correlates multiple potential problems into one concise indication of overall risk by weighing the DTT values for all metrics. For example, an element with multiple metrics that are each two days from breaching a threshold is ranked as a more severe risk than an element with just one metric that is only one day from breaching a threshold.

When a metric is rising, a DTT forecast is automatically produced. If the DTT is less than 1,000 days, it is recorded; if the DTT is more than 1,000 days, it is assigned a null value, which is displayed as N/A in the reports. If a metric has already exceeded a threshold value and is still rising, the DTT will be a negative number indicating the number of days past the threshold. If a metric is declining, whether it is currently over threshold or not, the DTT forecast is always null.

For the Frame Relay Forecast report, the forecast metrics are Input Utilization and Output Utilization. The Frame Relay ReportPack utilizes information collected from the RFC 1315 Frame Relay MIB. TREND performs a least squares linear regression analysis on data samples from the baseline period to forecast when either Utilization metric will reach the threshold. The default threshold values are predefined by TREND when the Frame Relay ReportPack is installed.

To provide a higher quality forecast for planning purposes, TREND bases the linear regression analysis on the 95th percentile of utilization samples taken during the day rather than on a simple average. The 95th percentile is the value below which 95% of the samples fall. The 95th percentile better characterizes the shape of the data by taking into account the variations or spikes in network utilization that occur during the day rather than smoothing the data out by averaging. By using 95th percentile instead of average utilization values, TREND forecasts provide better advanced warning and bandwidth figures that will handle the majority of spikes in network traffic.

A forecast that is based on simple average utilization does not adequately show when normal spikes will overload the network and cause lost productivity and lower service levels to end users.

Figure 3-7 shows a compressed sample baseline data set of only 3 weeks, which is used to illustrate the difference between forecasts based on 95th percentile and average utilization. The 95th percentile forecast clearly gives a better planning range for bandwidth that will handle more of the fluctuations in traffic demands.

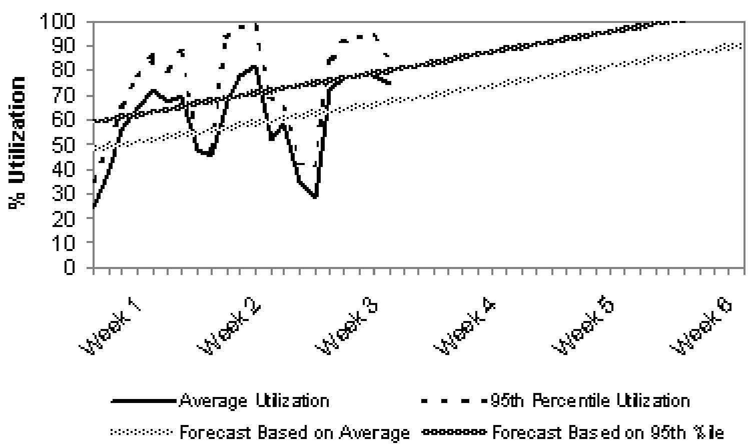

Figure 3-7: Forecasts Based on 95th Percentile vs. Average Utilization

To assist you in analyzing long-term trends, each Frame Relay PVC projected to exceed exception thresholds is linked to multiple drill-down components immediately below the Estimated Days to Threshold (DTT) table.

Grade of Service Drill-Down

Figure 3-8 shows the Grade of Service stacked bar chart, which presents a composite grade that correlates a number of important service quality metrics appropriate to the managed element or resource.

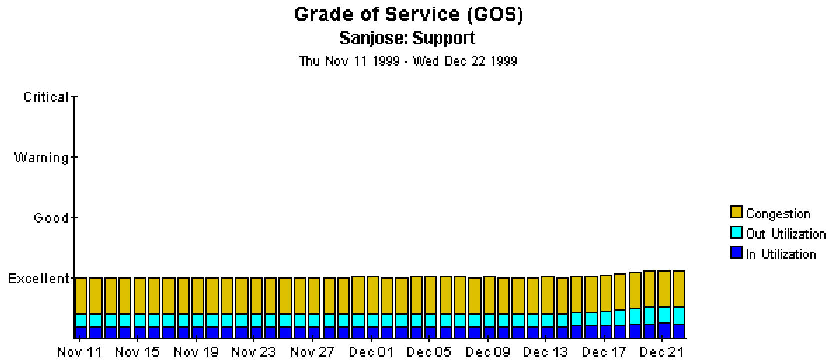

Figure 3-8: Forecast—Grade of Service Drill-Down (Frame Relay)

Health scores for each metric are assigned and weighted according to their relative importance, added together with other metric scores, and a final grade is automatically calculated. If the default polling interval is used, samples are taken every 15 minutes, so a daily Grade of Service (GOS) grade represents the average of the weighted composite scores for 96 daily samples. By transforming numerous service-related statistics into an easily understood stacked bar graph and presenting it over the baseline period, the GOS chart helps you recognize changes in Frame Relay PVC health and the root cause of those changes.

The Frame Relay GOS metrics are Input Utilization, Output Utilization, and Congestion. Utilization is expressed as a percent of CIR (Committed Information

Rate) if that value is available; otherwise, it is expressed as a percent of ifSpeed (interface speed). A health score is assigned to each metric. A value of 1 corresponds to an Excellent rating.

Since not all metrics may impact service quality equally, each metric is multiplied by a weighting factor. For Frame Relay, Input and Output Utilization each comprise 20% of the final GOS score, and Congestion comprises the remaining 60%.

By assigning a grade and corresponding description (for example, Excellent, Good, Warning, Critical) to each time interval, TREND provides an aggregate view of the stability or changes in overall health for the element.

Utilization Drill-Down

Figure 3-9 shows the Utilization line graph, which presents input and output utilization values both as a percent of CIR.

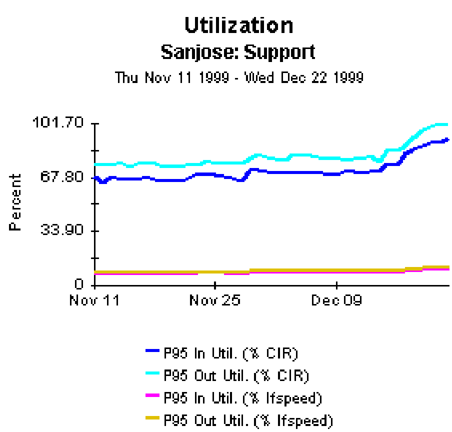

Figure 3-9: Forecast—Utilization Drill-Down (Frame Relay)

If a CIR has been defined, this graph helps identify PVCs with utilization exceeding the CIR, which can cause congestion and frames to be discarded. By reviewing the

changes and trends in utilization over the baseline period, you can immediately understand whether utilization that exceeds CIR represents is a brief period of heavy traffic that is not likely to be repeated or a long-standing situation of steadily growing traffic.

Congestion Drill-Down

Figure 3-10 shows the Congestion line graph, which itemizes the number of Forward Explicit Congestion Notifications (FECNs) and (Backward Explicit Congestion Notifications (BECNs) that have been received for that PVC. Congestion notifications are indications of degrading service to your end users.

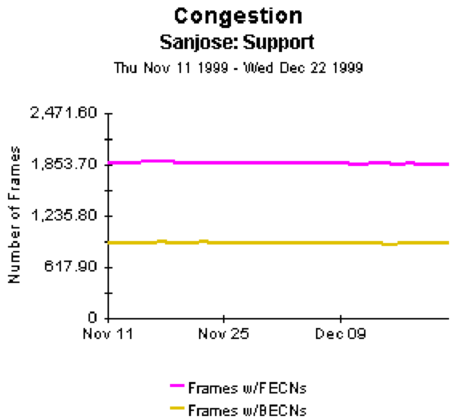

Figure 3-10: Forecast—Congestion Drill-Down (Frame Relay)

Capacity Planning Report

The Capacity Planning report details the most overutilized and the most underutilized interfaces. This information can be used to:

- Point out opportunities for balancing the load more evenly between interfaces to improve service levels without requiring additional investment.
- Identify interfaces where more bandwidth will be required to meet user demands.
- Expose interfaces where a lower, less-costly service would still meet user needs.

For example, the following components comprise the Capacity Planning report for WAN Connectivity ReportPack:

- Overutilized WAN Interfaces Summary summary table (see "Overutilized WAN Interfaces Summary" on page 3-14) and the following drill-down components:
    - Daily GOS Drill-Down Chart (see "Daily Grade of Service Drill-Down" on page 3-15).
    - Daily Utilization Drill-Down Chart (see "Daily Utilization Drill-Down" on page 3-16).
    - Daily Errors and Discards Drill-Down Chart (see "Daily Errors & Discards Drill-Down" on page 3-17).
- Underutilized WAN Interfaces Summary table and the following drill-down components:
    - Daily GOS Drill-Down Chart
    - Daily Utilization Drill-Down Chart
    - Daily Errors and Discards Drill-Down Chart

Note: **The following sections describe overutilized interfaces only, as the analysis information for underutilized interfaces is the same.**

The following sections describe these charts in more detail. For more information about the WAN Connectivity ReportPack, see “LAN and WAN ReportPacks” on page 5-1.

Overutilized WAN Interfaces Summary

Figure 3-11 shows the Overutilized WAN Interfaces summary table, which lists the top 10 interfaces projected to exceed the utilization threshold in the next 90 days.

Overutilized WAN Interfaces

Projected To Exceed 80% Utilization Within 90 Days

All Values Based On 95th Percentile (P95)

Device	Element	Description	Current P95 In Util	Projected In Util 30 Days	Projected In Util 60 Days	Projected In Util 90 Days	Current P95 Out Util
lincoln	1	"WAN Link: Lincoln-Tulsa"	93.15	100.98	108.69	116.41	90.84
boston	3	"WAN Link: Boston-San Jose "	86.09	94.82	102.63	110.44	87.08
sanjose	3	"WAN Link: New York"	55.48	66.94	77.43	87.91	66.51
miami	3	"WAN Link: New York"	55.47	67.25	77.70	88.14	67.34
sanluis	1	"WAN Link: Atlanta-Portland"	84.57	81.87	80.91	79.96	85.38
tulsa	3	"WAN Link: Tulsa-Lincoln -II"	85.68	82.56	81.87	81.18	85.87

Figure 3-11: Capacity Planning—Overutilized WAN Interfaces Summary

By comparing this information with the interfaces listed in the Underutilized WAN Interfaces Summary table, you can potentially improve performance by shifting the traffic load from overutilized to underutilized links. If that is not possible, the Overutilized WAN Interfaces Summary table provides advanced warning of interfaces where more bandwidth will soon be needed to prevent congestion from affecting end user service levels.

The Overutilized WAN Interfaces Summary table displays average Input and Output Utilization values for the baseline period for each interface, and specific forecast

values for 30, 60, and 90 days into the future. The forecast values are the result of a least squares linear regression analysis performed on data samples from the baseline period. (TREND uses a rolling baseline. For an explanation of how baselines are calculated, see “Baselines” on page 2-10.)

The overutilized interfaces are sorted from the highest forecast 90-day utilization (either input or output) value to lowest. Only interfaces projected to exceed the utilization threshold within 90 days are included in the table. Depending on your network performance profile, none of the WAN interfaces may meet this criteria or fewer than 10 interfaces may be listed. To assist you in analyzing long-term capacity trends, each overutilized interface is linked to multiple drill-down components immediately below the table.

Daily Grade of Service Drill-Down

Figure 3-12 shows the Daily Grade of Service stacked bar chart, which presents a composite grade that correlates a number of important service quality metrics appropriate to the managed element or resource.

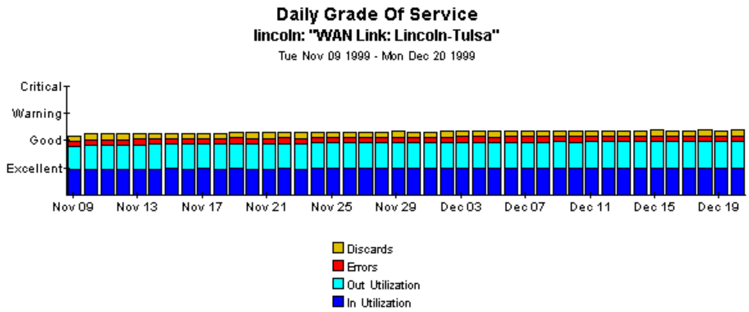

Figure 3-12: Capacity Planning—Daily Grade of Service Drill-Down (WAN)

The Health scores for each metric are assigned and weighted according to their relative importance, added to other metric scores, and a final grade is calculated. If samples are taken every 15 minutes, a daily Grade of Service (GOS) grade represents the average of the weighted composite scores for 96 daily samples. By

transforming numerous service-related statistics into an easily understood stacked bar graph and presenting it over the baseline period, the GOS chart helps you recognize changes in WAN interface health and the root cause of those changes.

The WAN GOS metrics are Input Utilization, Output Utilization, Errors, and Discards. A health score is assigned to each metric. A value of 1 corresponds to an Excellent rating.

Daily Utilization Drill-Down

Figure 3-13 shows the Daily Utilization line graph for overutilized interfaces, which allows you to review the changes and trends in utilization over the baseline period, so you can immediately understand whether the overutilized interfaces represent a brief period of heavy traffic or a long-standing situation of steadily growing traffic. Since both In and Out Utilization in Figure 3-13 increase over time, you should consider a PVC upgrade in the future.

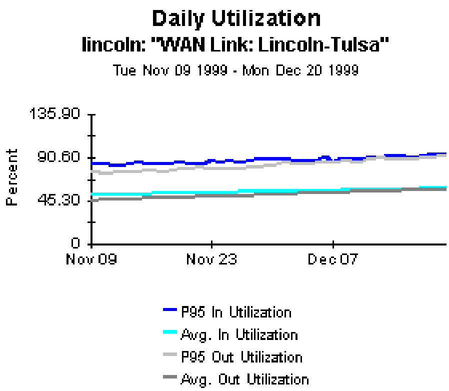

Figure 3-13: Capacity Planning—Daily Utilization Drill-Down (WAN)

Daily Errors & Discards Drill-Down

Figure 3-14 shows the Daily Errors and Discards line graph for overutilized interfaces. The ratio of errors to total packets (inErrors:inPackets and outErrors:outPackets) is calculated for each polling cycle. A similar calculation is done to reveal the ratio of discards to total packets. The chart shows the daily 95th percentile of these calculations as a percentage of total packets to assist you in determining whether a sudden spike has occurred or a long-standing, low-level problem has finally surfaced. This example shows that the ratio of errors and discards to total packets is fairly stable over time.

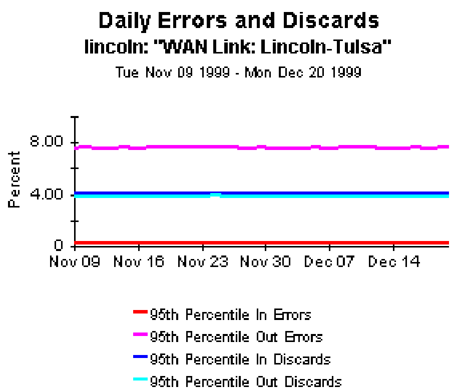

Figure 3-14: Capacity Planning—Daily Errors and Discards Drill-Down (WAN)

Top Ten

The Top Ten report identifies the top ten elements with the highest ranked metric values. The top ten elements are listed by highest rank and highest change of rank.

For example, the Top Ten report for the Bay Router ReportPack lets the network manager know at-a-glance which Bay routers produce the greatest throughput and

exhibit the poorest health. The top ten Bay routers, ranked by total throughput (packets), are listed by highest rank as well as highest change of rank. Throughput and throughput-change charts alert you to the routers' greatest changes in throughput, suggesting potential capacity overloading if increasing trends continue. The throughput change report also may highlight the newly installed devices and technologies that are approaching a logical utilization threshold or are underutilized.

The following components comprise the Top Ten report for Bay Router:

- Busiest Bay Routers summary table (see "Busiest Bay Routers" on page 3-18).
- Largest Bay Routers Throughput Changes summary table (see "Largest Bay Router Throughput Changes" on page 3-19).
- Bay Routers with Most Percentage of Discards summary table (see "Bay Routers with the Most Percentage of Discards" on page 3-19).
- Bay Routers with Greatest Change in Percentage of Discards summary table (see "Bay Routers with the Greatest Change in Percentage of Discards" on page 3-20).

The following sections describe these charts in more detail. For detailed information about the Bay Router ReportPack, see "Bay Router ReportPack" on page 7-1.

Busiest Bay Routers

Figure 3-15 shows the **Busiest Bay Routers** summary table, which lists up to ten Bay routers with the highest throughput (packets) for the report day.

Busiest Bay Routers

Sun Dec 19 1999

Router	Pct Discards	Pkts/sec	Rank	Bytes/sec	Prev Rank
sanjose	3.73	27538522.93	4	1182195575.79	6
sanluis	3.48	27366425.47	6	853449719.02	5

Figure 3-15: Top Ten—Busiest Bay Routers

Largest Bay Router Throughput Changes

Figure 3-16 shows the **Largest Bay Router Throughput Changes** summary table, which ranks the selected routers in ascending sequence by change in throughput (packets).

Largest Bay Routers Throughput Changes

Sun Dec 19 1999

Router	Pct Discards	Pkts/sec	Rank	Change Rank	Prev Pkts/sec	Prev Rank
sanjose	3.73	27538522.93	4	1	27339710.09	6
sanluis	3.48	27366425.47	6	7	27378710.53	5

Figure 3-16: Top Ten—Largest Bay Router Throughput Changes

The report shows the current day's discards and traffic (in packets-per-second); previous day's traffic (in packets-per-second); and current day's rank, previous day's rank, and change in rank from the previous day. A large change in rank identifies a volatile resource, and you should investigate the reasons for the high change in throughput.

Bay Routers with the Most Percentage of Discards

Figure 3-17 shows the **Bay Routers with the Most Percentage of Discards** summary table, which ranks Bay routers in descending order by the highest percentage of packet discards for the report day. Packets/second and bytes/second for the selected Bay routers are also shown.

Bay Routers with Most Percentage of Discards

Sun Dec 19 1999

Device	Pct Pkts Discarded	Pkts/sec	Rank	Bytes/sec	Prior Rank
sanjose	3.73	27538522.93	3	1182195575.79	3
sanluis	3.48	27366425.47	4	853449719.02	4

Figure 3-17: Top Ten—Bay Routers with Most Percentage of Discards

3 Analyzing TREND Reports

Bay Routers with the Greatest Change in Percentage of Discards

Figure 3-18 shows the **Bay Routers with the Greatest Change in Percentage of Discards** summary table, which ranks the selected routers in ascending sequence by change in discard percentage from the previous day. A large change in rank identifies a volatile resource, and you should investigate the reasons for the high change in discard percentage.

Bay Routers With Greatest Change in Percentage of Discards

Wed Jan 05 2000

Device	Pct Pkts Discarded	Pkts/sec	Rank	Bytes/sec	Change Rank	Previous Pct Pkts Discarded	Prior Rank
sanluis	3.48	27366425.47	4	853449719.02	4	3.48	4
sanjose	3.73	27538522.93	3	1182195575.79	6	3.73	3

Figure 3-18: Top Ten—Bay Routers with the Greatest Change in Percentage of Discards

QuickView and Snapshot Reports

The QuickView report identifies the ten elements with the highest (worst) grade for scored metrics. Network managers and analysts can go to this report to select an element and assess its status in several performance areas.

Alternatively, the Snapshot report shows the same information, however, you select elements you want to view from a pick list when you invoke the report. For more information about selecting an element for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

For example, the QuickView and Snapshot reports for the Frame Relay ReportPack identify the top ten PVCs with the highest (worst) grade of service scores (according to Grade of Service score). Network managers and analysts can go to this report to

select a PVC from a grade of service chart, and assess its status in several performance areas. Drill-down charts show daily distribution of bandwidth, as polled volume against the baseline and congestion at the as polled level.

The following sections describe these charts in more detail. For detailed information about the Frame Relay ReportPack, see "Frame Relay ReportPack" on page 4-1.

The QuickView and Snapshot reports comprise the following charts:

- (QuickView) **Daily Grade of Service** chart (see "Daily Grade of Service (QuickView)" on page 3-22).
- (Snapshot) **Snapshot Selection** table (see "Snapshot Selection Table (Snapshot)" on page 3-23).
    - (Snapshot) The **Grade of Service** chart (see "Grade of Service (Snapshot)" on page 3-23).
    - The **Hourly Volume** graph (see "Hourly Volume" on page 3-24).
    - The **Hourly Congestion** graph (see "Hourly Congestion" on page 3-24).
    - The **Distribution of Bandwidth Utilization - Input** and **Distribution of Bandwidth Utilization - Output** charts (see "Distribution of Bandwidth Utilization - Input" on page 3-25 and "Distribution of Bandwidth Utilization - Output" on page 3-26).

Daily Grade of Service (QuickView)

Figure 3-19 shows the top ten PVCs with the highest daily GOS scores. Table 4-3 describes how Grade of Service is calculated.

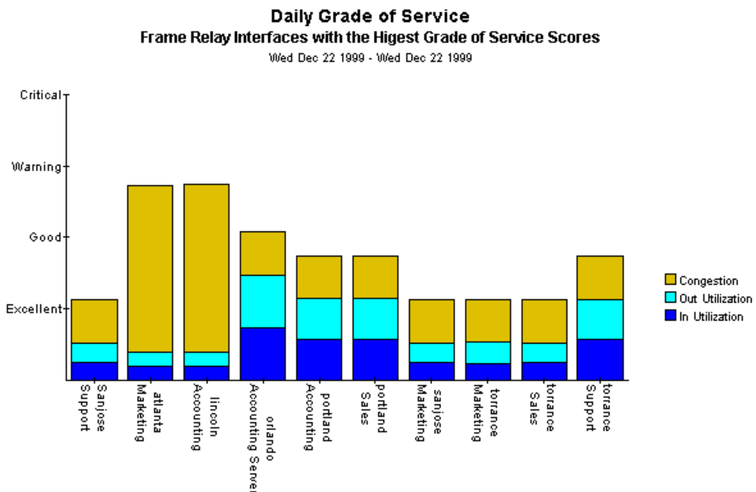

Figure 3-19: QuickView—Daily Grade of Service (Frame Relay)

Double-click on a device in the chart to update the content of the associated drill-down components.

Snapshot Selection Table (Snapshot)

Figure 3-20 shows the **Snapshot Selection** table, which lists the time period for which the PVC information was collected, the PVC name, the interface, the PVC's Input Utilization (CIR and ifSpeed), Output Utilization (CIR and ifSpeed), and Volume.

Snapshot Selection

Mon Dec 27 1999

Device	Interface/DLCI	Description	Congestion	Input Utilization (CIR)	Input Utilization (IfSpeed)	Output Utilizatior (CIR)
Sanjose	1.20		0.04	133.29	N/A	149.49
atlanta	1.24		0.82	0.00	0.00	0.00

Figure 3-20: Snapshot—Snapshot Selection Table (Frame Relay)

Double-click on a device in the table to update the content of the associated drill-down components with performance data for the device.

Grade of Service (Snapshot)

Figure 3-21 shows the Grade of Service stacked bar chart, which displays the selected PVC's GOS scores for each sample taken over the report day.

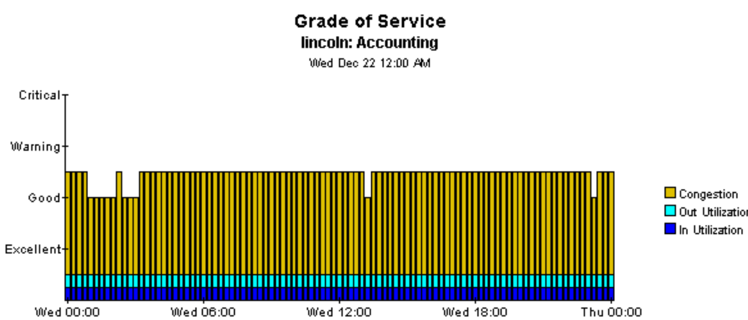

Figure 3-21: Snapshot—Grade of Service (Frame Relay)

Hourly Volume

Figure 3-22 shows the Hourly Volume area graph, which displays the total hourly volume for the selected PVC over the report day.

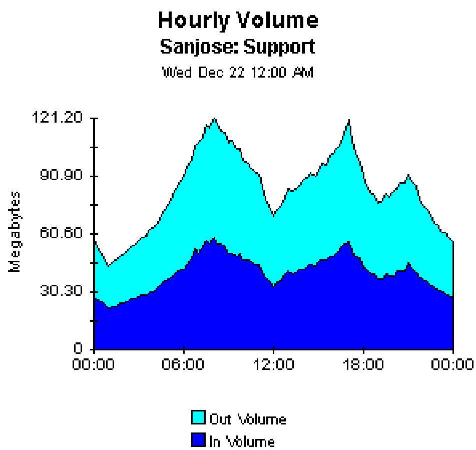

Figure 3-22: QuickView—Hourly Volume (Frame Relay)

Hourly Congestion

Figure 3-23 shows the Hourly Congestion line graph, which plots the total number of frames for each hour with Forward Explicit Congestion Notification and the total

number of frames with Backward Explicit Congestion Notification for the selected PVC for each sample taken during the report day.

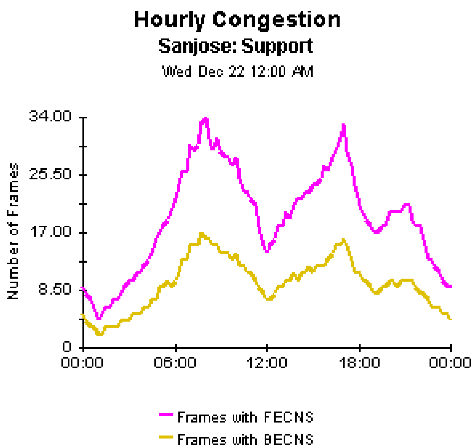

Figure 3-23: QuickView—Hourly Congestion

Distribution of Bandwidth Utilization - Input

Figure 3-24 shows the Distribution of Bandwidth Utilization - Input pie chart, which shows the percentage of samples taken during the day in which the In Utilization

metric calculated for the samples fall in a defined range. For more information about frequency distribution, see "Frequency Distribution" on page 2-14.

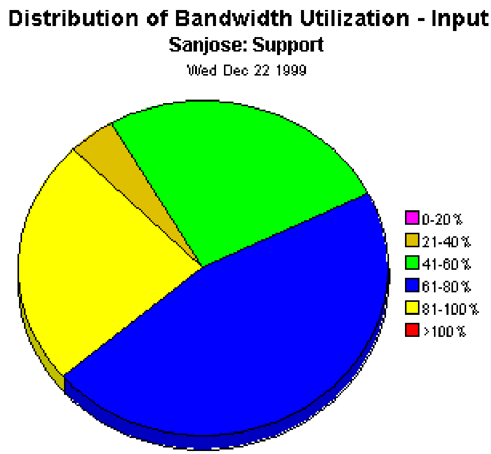

Figure 3-24: QuickView—Distribution of Bandwidth-Input

Distribution of Bandwidth Utilization - Output

Figure 3-24 show Distribution of Bandwidth Utilization-Output pie chart, which displays the percentage of samples taken during the day in which the Out Utilization

metric calculated for the samples fall in a defined range. For more information about frequency distribution, see "Frequency Distribution" on page 2-14.

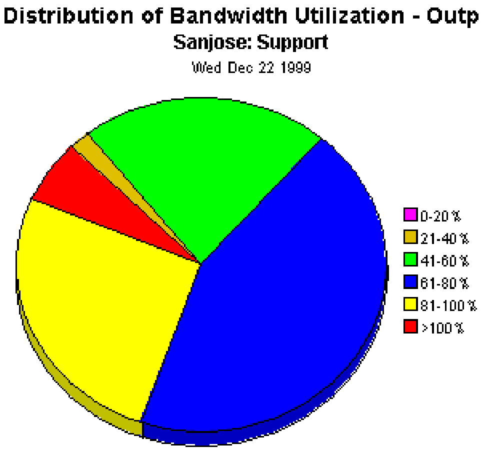

Figure 3-25: QuickView—Distribution of Bandwidth-Output

Near Real Time QuickView and Snapshot Reports

The Near Real Time--QuickView and the Near Real Time--Snapshot reports provide device performance statistics up to the last SNMP poll. The reports does not rely on nightly summaries; therefore, they provide instant reporting on collected data. The QuickView lets you select a device from a table, and assess its status in several performance areas.

Alternatively, the Snapshot report shows the same information as the QuickView report, however, you select the device you want to view from a pick list when you invoke the Snapshot report. For more information about selecting a device for a Snapshot report, see "Selecting Deferred Targets or Elements for a Snapshot Report" on page 1-27.

The following ReportPacks provide Near Real Time QuickView and Snapshot reports:

- Bay Router
- Cisco Router
- Frame Relay
- LAN/WAN
- System

For more information about the Near Real Time reports for each ReportPack, see the appropriate chapter for the ReportPack in this guide.

For example, the Near Real Time QuickView and Near Real Time Snapshot reports for the Bay Router ReportPack comprises the following table and graphs:

- Router Selection List table (see “Router Selection List” on page 3-29).

  Double-click on a router in the table to update one of the following drill-down charts with performance data for the selected router:

  - CPU Utilization chart (see “CPU Utilization Drill-Down” on page 3-30).
  - Memory Utilization chart (see “Memory Utilization Drill-Down” on page 3-31).
  - Buffer Utilization chart (see “Buffer Utilization Drill-Down” on page 3-32).

Router Selection List

Figure 3-26 shows each router and its performance statistics based on the data polled for the last 6 hours. The statistic shown for each metric is the average of all the data collected for that statistic over the 6-hour period.

Router Selection List

Select router to see near real time reports

Router	Slot	CPU Utilization	Memory Utilization	Buffer Utilization
sanluis	1	50.32	46.81	46.84
sanjose	3	50.30	47.27	46.92
sanjose	5	50.03	46.01	46.63
sanjose	1	37.88	37.82	37.67
sanluis	3	33.38	35.03	33.04

Figure 3-26: Near Real Time—Router Selection List (Bay Router)

Double-click on a device in the table to update the content of the associated drill-down components with performance data for the device.

CPU Utilization Drill-Down

Figure 3-27 shows the **CPU Utilization** line graph, which displays the CPU utilization for each sample taken during the report period for the selected router.

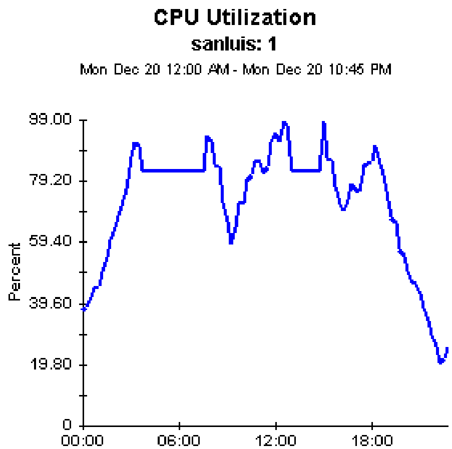

Figure 3-27: Near Real Time—CPU Utilization Drill-Down (Bay Router)

Memory Utilization Drill-Down

Figure 3-28 shows the Memory Utilization line graph, which displays memory utilization for each sample taken during the report period for the selected router.

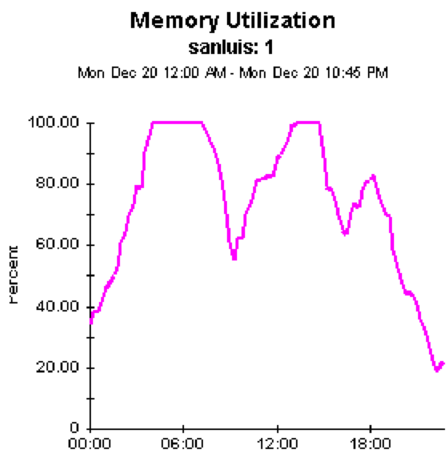

Figure 3-28: Near Real Time—Memory Utilization Drill-Down (Bay Router)

3 Analyzing TREND Reports

Buffer Utilization Drill-Down

Figure 3-29 shows the **Buffer Utilization** line graph, which displays the buffer utilization for each sample taken during the report period for the selected router.

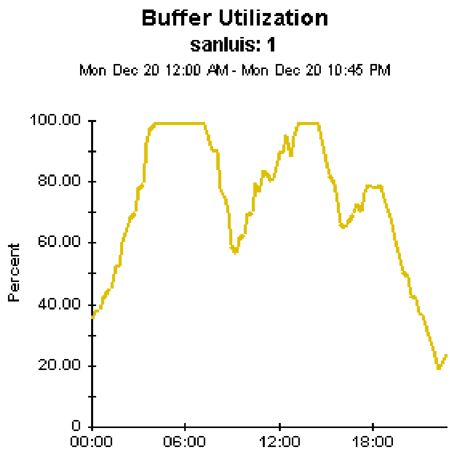

Figure 3-29: Near Real Time—Buffer Utilization Drill-Down (Bay Router)

Hot Spots Report

The Hot Spots report identifies areas where threshold values have been exceeded (for example, excessive utilization or error counts). Intelligent exception thresholds are predefined in TREND, but you can customize them to suit your enterprise.

For example, the following components comprise the Hot Spots report for the Cisco Router ReportPack:

- Problem Summary for the Day table (see "Problem Summary for the Day" on page 3-33) and the following drill-down components:

- Hourly Grade of Service chart (see "Hourly Grade of Service Drill-Down" on page 3-34).
- CPU Utilization chart (see "CPU Utilization Drill-Down" on page 3-36).
- Buffer Utilization chart (see "Buffer Utilization Drill-Down" on page 3-37).
- Memory Utilization chart (see "Memory Utilization Drill-Down" on page 3-38).
- Percent Discards chart (see "Percent Discards Drill-Down" on page 3-39).
- Buffer Faults chart (see "Buffer Faults Drill-Down" on page 3-39).
- Exception Detail summary table (see "Exception Detail Drill-Down" on page 3-40).

The following sections describe these charts in more detail. For more information about the Cisco Router ReportPack, see "Cisco Router ReportPack" on page 8-1.

Problem Summary for the Day

Figure 3-30 shows the Problem Summary for a Day table, which summarizes multiple exception conditions into one aggregate total number (Total Exceptions) indicating the top ten offenders, ranked by total number of exceptions, over the day.

Problem Summary for the Day
Cisco Routers With Highest Exception Counts
Sun Dec 19 1999

Router	CPU Utilization Exceptions	Memory Utilization Exceptions	Buffer Utilization Exceptions	Buffer Failure Exceptions	Total Exceptions
houston	47.00	48.00	0.00	148.00	243.00
cotati	0.00	45.00	96.00	0.00	141.00
tulsa	7.00	0.00	0.00	0.00	7.00
boston	7.00	0.00	0.00	0.00	7.00

Figure 3-30: Hot Spots— Problem Summary for the Day (Cisco Router)

To help you analyze router problems, the summary row for each device is linked to multiple drill-down components that bring together the varied information required to investigate potential network trouble spots, pinpoint the source of the difficulty, and determine possible resolutions.

For Cisco routers, exception conditions are tracked for CPU Utilization, Memory Utilization, Buffer Utilization, and Buffer Faults. Reporting these metrics helps you identify routers that are unable to handle the amount of network traffic being generated. These indicators tell you when a router needs a CPU upgrade, when there is inadequate memory, or when buffers are being overutilized.

The drill-down components below the Daily Problem Summary chart supply instant access to detailed information in the areas of Grade of Service, CPU Utilization, Buffer Utilization, Memory Utilization, Discards, and Buffer Faults for each hour of the last day. You can readily see the correlation between critical performance indicators during particular time periods. The charts and graphs show exactly when a threshold is breached and what happens before and after the breach. Thus, you can determine if a problem is a long-standing issue that has been growing steadily or a sudden burst that may not be repeated.

Hourly Grade of Service Drill-Down

Figure 3-31 shows the Hourly Grade of Service stacked bar chart, which presents a composite grade that correlates a number of important service quality metrics appropriate to that element or resource.

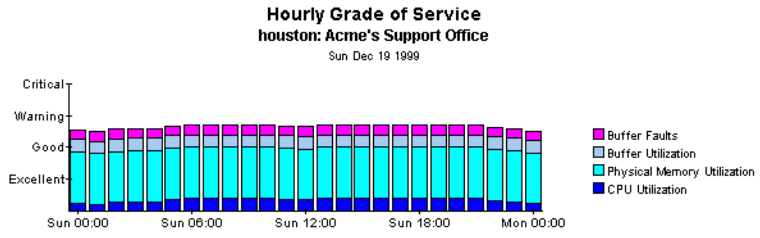

Figure 3-31: Hot Spots—Hourly Grade of Service Drill-Down (Cisco Router)

Health scores for each metric are assigned and weighted according to their relative importance, added to other metric scores, and a final grade is calculated. By transforming numerous service-related statistics into an easily understood stacked bar graph and presenting it over the critical troubleshooting period of the last 24 hours, this graph helps you recognize changes in Cisco router health and the root cause of those changes.

The Grade of Service (GOS) metrics for Cisco routers are CPU Utilization, Physical Memory Utilization, Buffer Utilization, and Buffer Faults. A health score is assigned to each metric. A value of 1 corresponds to an Excellent rating. Because not all metrics may impact service quality equally, each metric is multiplied by a weighting factor. By assigning a grade and corresponding description (for example, Excellent, Good, Warning, Critical) to each time interval, TREND provides an aggregate view of the stability or changes in overall health for the element.

In Figure 3-31, router health starts at Good and degrades to Warning at 5:00 a.m., largely due to increased CPU and physical memory utilization. The CPU Utilization drill-down chart (Figure 3-32) provides details on the exact changes that took place. You can further investigate buffer utilization in the Buffer Utilization drill-down chart (Figure 3-33). The physical memory utilization (Figure 3-34) is consistently high. This indicates a memory upgrade is required. The Buffer Faults drill-down chart (Figure 3-36) reveals a minimal number of faults; the faults that do exist correlate to the CPU Utilization increases. Together, these charts indicate a router that is pushing the limits of its CPU and physical memory resources. It is not yet discarding much traffic so end-user service levels are not severely impacted, but the need for an upgrade is indicated.

CPU Utilization Drill-Down

Figure 3-32 shows the CPU Utilization line graph shows hour-by-hour changes and trends in CPU utilization for the selected Cisco router.

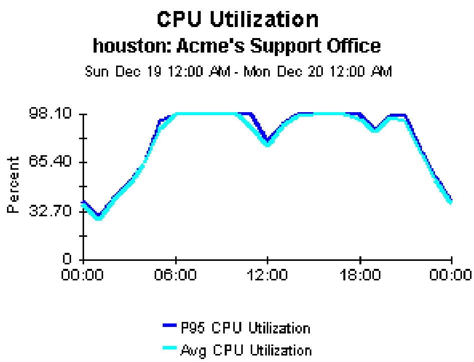

Figure 3-32: Hot Spots—CPU Utilization Drill-Down (Cisco Router)

By showing the time of day when an exception occurs, the drill-down components can help narrow possible causes of the problem. For example, a consistent spike in utilization at 11:00 a.m. might indicate that a resource-intensive application (such as a backup) is being run. In that case, you can schedule the application to run at some other time, such as 11:00 p.m., to balance the router load. Figure 3-32 shows CPU utilization is following the normal workday hours

Buffer Utilization Drill-Down

Figure 3-33 shows the Buffer Utilization line graph, which displays the percentage of the various size buffers (Huge Buffer, Big Buffer, Large Buffer, Medium Buffer, and Small Buffer) that are in use during each time interval over the past 24 hours.

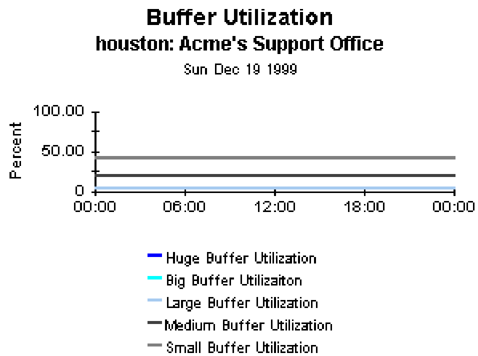

Figure 3-33: Hot Spots—Buffer Utilization Drill-Down (Cisco Router)

Memory Utilization Drill-Down

Figure 3-34 shows the Memory Utilization line graph. Memory utilization is consistently high, indicating that a memory upgrade is required.

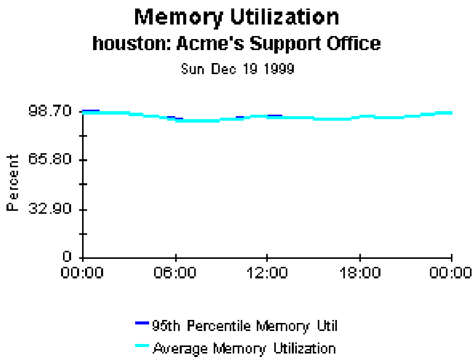

Figure 3-34: Hot Spots—Memory Utilization Drill-Down (Cisco Router)

Percent Discards Drill-Down

Figure 3-35 shows the Percent Discards bar chart for each hour of the day. If a high discard percentage correlates with a high CPU utilization number, the router is having trouble keeping up with the traffic volume.

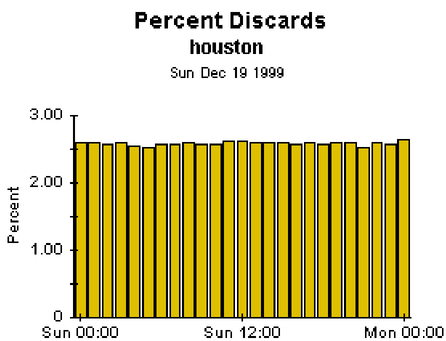

Figure 3-35: Hot Spots—Percent Discards Drill-Down (Cisco Router)

Buffer Faults Drill-Down

Figure 3-36 shows the Buffer Faults bar chart, which displays the total number of buffer faults for each hour of the day. If a high number of buffer faults correlates

with times of high CPU utilization, the router is having trouble keeping up with the traffic volume.

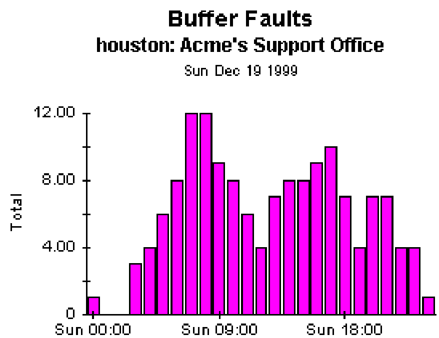

Figure 3-36: Hot Spots—Buffer Faults Drill-Down (Cisco Router)

Exception Detail Drill-Down

Figure 3-37 shows the Exception Detail table, which lists a router's CPU Utilization, Memory Utilization, Buffer Utilization, and Buffer Faults for each sample taken over the report day.

Exception Detail
houston: Acme's Support Office
Sun Dec 19 12:00 AM

Time Period	CPU Util	Mem Util	Buf Util	Buf Flts
Sun Dec 19 12:00 AM	45.00	97.63	22.90	1.00
Sun Dec 19 12:15 AM	40.00	97.81	22.90	0.00
Sun Dec 19 12:30 AM	34.00	98.17	22.90	0.00
Sun Dec 19 12:45 AM	31.00	98.47	22.90	0.00
Sun Dec 19 01:00 AM	23.00	98.76	22.90	0.00
Sun Dec 19 01:15 AM	26.00	98.65	22.90	0.00
Sun Dec 19 01:30 AM	29.00	98.45	22.90	0.00

Figure 3-37: Hot Spots—Exception Detail (Cisco Router)

Service Level Management Report

The Service Level Management Report provides an instant view of latency and availability plotted against contracted service levels. Managers and customers can see at-a-glance if service levels have been achieved for the previous day, month, or other time period.

For example, the System ReportPack is based on Empire Technology's SystemEDGE™ product and utilizes data collected from MIB-II.

The following components comprise the Service Level Management report for the System ReportPack:

- Availability table (see "Availability Table" on page 3-42) and the following drill-down component:
  - Daily Availability chart (see "Daily Availability Drill-Down" on page 3-42).
  - Hourly Availability chart (see "Hourly Availability Drill-Down" on page 3-43).
- Response Time table (see "Response Time" on page 3-45) and the following drill-down components:
  - Daily Response Time chart (see "Daily Response Time Drill-Down" on page 3-46).
  - Hourly Response Time chart (see "Hourly Response Time Drill-Down" on page 3-47).

The following sections describe these charts in more detail. For more information about the System ReportPack, see "System ReportPack" on page 12-1.

Availability Table

Figure 3-38 shows the Availability table and its Daily and Hourly Availability drill-down charts. The Availability table lists the systems with the lowest availability.

TREND processes each system's availability results and highlights problem systems whose availability is below contracted Service Level Agreement (SLA) metrics. TREND implements a management-by-exception policy by sifting through all managed systems and focusing your attention on only the problem units. This reduces the amount of data you must review to stayed informed on system health.

Availability
Systems With Lowest Availability
Thu Oct 21 1999

Device	Availability
wasp	110.94
leyte	110.94
hornet	110.94
laffey	110.94
nimitz	110.94
yorktown	110.94
eisenhower	110.94
enterprise	110.94

Figure 3-38: Service Level—Availability Table (System)

When the System ReportPack is installed, the Availability SLA metric is set to 99.5% by default. The Availability chart shows the 10 systems with the lowest availability.

Daily Availability Drill-Down

Figure 3-39 shows the Daily Availability line graph. You can double-click any system in the Availability table to update this graph with the average daily

availability over the report period (previous 42 days) for analysis and investigation of availability changes.

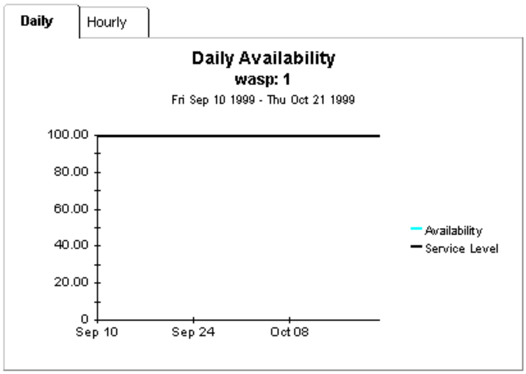

Figure 3-39: Service Level—Daily Availability Drill-Down (System)

Availability is a critical metric in most SLAs. In principle, availability is simply the percent of time that the system is up during any given time period. Accessible data to measure availability and the quality of that data differ from device to device.

Hourly Availability Drill-Down

Figure 3-40 shows the Hourly Availability line graph. You can double-click any system in the Availability table to update this graph with the average hourly

availability over the report period (previous day) for analysis and investigation of availability changes.

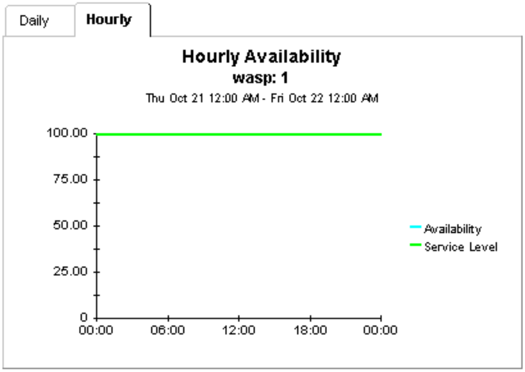

Figure 3-40: Service Level—Hourly Availability Drill-Down (System)

Availability is a critical metric in most SLAs. In principle, availability is simply the percent of time that the system is up during any given time period. Accessible data to measure availability and the quality of that data differ from device to device.

Response Time

Figure 3-41 shows the Response Time table, which lists the systems with the highest response times.

Response Time

Systems With Highest Respose Time

Thu Oct 21 1999

Device	Response Time
wasp	100.00
leyte	100.00
hornet	100.00
laffey	100.00
nimitz	100.00
yorktown	100.00
eisenhower	100.00
enterprise	100.00

Figure 3-41: Service Level—Response Time Drill-Down (System)

Response Time is another critical metric in most SLAs. TREND processes each system's response time results and highlights problem systems with the worst or longest response times in this chart. When vendor-specific response time data from our partners is not available, DeskTalk uses the time it takes for a device to respond to TREND's SNMP GET requests as a viable indicator of response time. Because this measurement includes the desired device response time, plus the processing time of the device's SNMP agent, plus the round trip network delay, you can gain more insight by reviewing response time trends or changes rather than focusing on the raw response time value. The Service Level Management report assists in this by presenting daily response time over the baseline period (see Figure 3-42) for long-term trend analysis and hourly response time over the past day (see Figure 3-42) for detailed investigation so you can see the relative changes in response time. Double-click any system in the Response Time table to reveal the detailed daily and hourly response time drill-down charts.

Daily Response Time Drill-Down

Figure 3-42 shows the Daily Response Time line graph, which shows the network response time over the baseline period for the selected system.

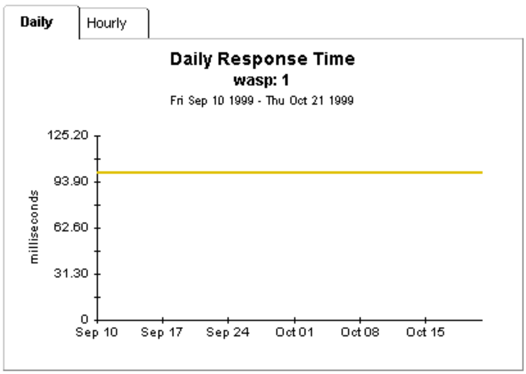

Figure 3-42: Service Level—Daily Response Time Drill-Down (System)

Hourly Response Time Drill-Down

Figure 3-43 shows the Hourly Response Time chart, which shows the network response time for the selected system for each hour yesterday.

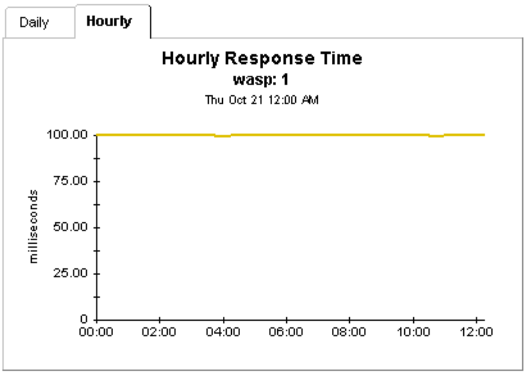

Figure 3-43: Service Level—Hourly Response Time Drill-Down (System)

TREND

4 Frame Relay ReportPack

The Frame Relay ReportPack provides comprehensive data collection and reporting to support the management of Frame Relay networks. It includes compound reports for capacity planning, service level reporting, forecasting, troubleshooting, and also provides an executive summary of the performance of all Permanent Virtual Circuits (PVCs) and reports the detail performance metrics for PVC networks.

The Frame Relay ReportPack answers questions such as:

- Are the Frame Relay networks providing an acceptable level of service to the customers, end-users, or both?
- What PVCs are overloaded and are there underutilized PVCs that can help balance the load?
- Are performance problems due to high levels of congestion, errors, excessive utilization, or other causes?
- What are the top ten PVCs contributing to the poor health of the network?
- Which PVCs are likely to degrade service in the near future if preventive action is not taken?

The Frame Relay ReportPack collects data from the RFC 1315 MIB and some proprietary MIBs, which are supported by most routers, Channel Service Units (CSUs)/

Data Service Units (DSUs), and other network devices. For more information about the Frame Relay ReportPack data source, see "Data Source" on page 4-25.

Report Descriptions

The following sections describe each report in the Frame Relay ReportPack in detail.

Executive Summary Report

The Executive Summary report provides executives and high-level managers with graphs that aggregate key metrics for all elements within a class. They can review the daily volume summary, and if desired, examine more detailed information about volume, grade of service, utilization, congestion, and total exception count.

Note: The TRENDweb Administrator can configure a Frame Relay ReportPack installation, so that the Executive Summary report will only display data for devices belonging to a specific customer. For more information about this feature, see "Displaying Data for Specific Customer Devices in Reports" on page 1-20**.**

The Executive Summary report comprises the following charts:

- The **Daily Volume for PVCs** bar chart shows the total Frame Relay volume over the baseline period (by default, the six weeks prior to and including the

report date), giving the executive an instant view of increases or decrease in Frame Relay usage over time.

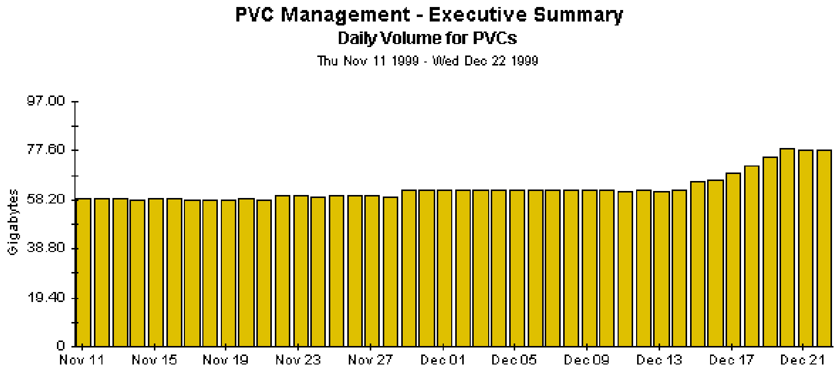

- The **Hourly Volume** bar chart shows the total volume for all PVCs for each hour of the report day.

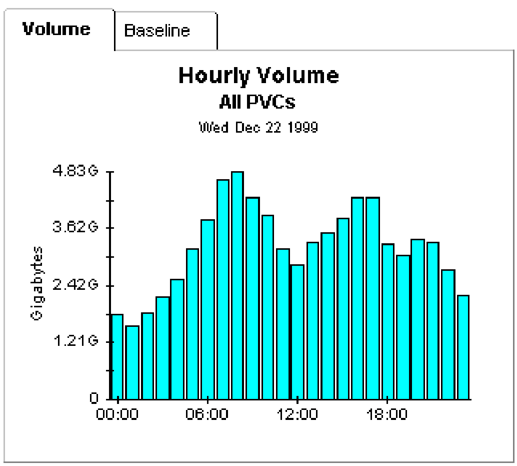

- The **Hour of Day Baseline** bar chart shows the average volume for each hour of the day over the baseline period. You can compare total volume for each hour of the report day with the baseline for that hour for the report period.

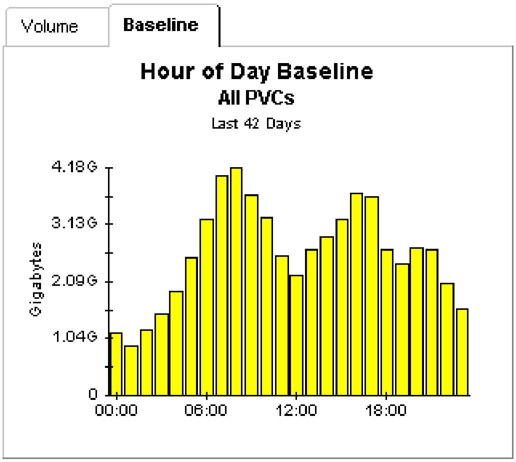

- The **Hourly Grade of Service** stacked bar chart provides an overview of PVC performance, providing weighted values for utilization and congestion. Table 4-3 describes how Grade of Service is calculated.

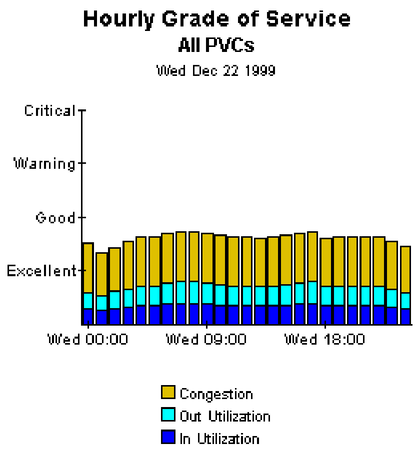

- The **Hourly Utilization** line graph calculates the average In Utilization and Out Utilization calculated two ways (one using CIR and one using ifSpeed) of all samples collected for each hour of the report day.

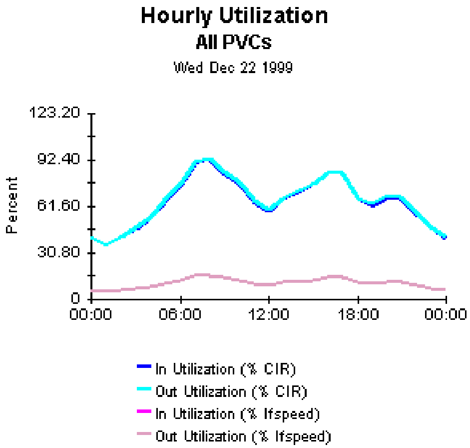

- The **Hourly Congestion** line graph shows total frames with Backward Explicit Congestion Notification (BECN) and Forward Explicit Congestion Notification (FECN) for all PVCs over each hour of the report day.

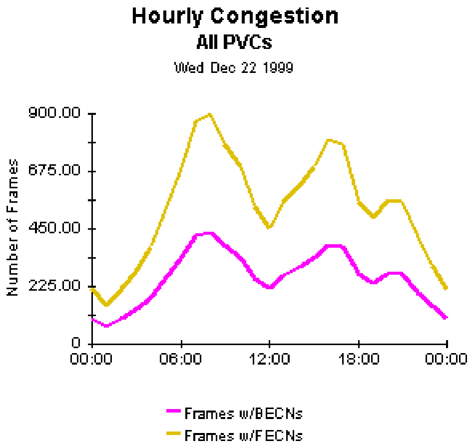

- The **Hourly Exception Count** stacked bar chart shows at a glance how many PVCs experience exception conditions for utilization and congestion over the day. It also includes individual plots of Utilization and Congestion over the day.

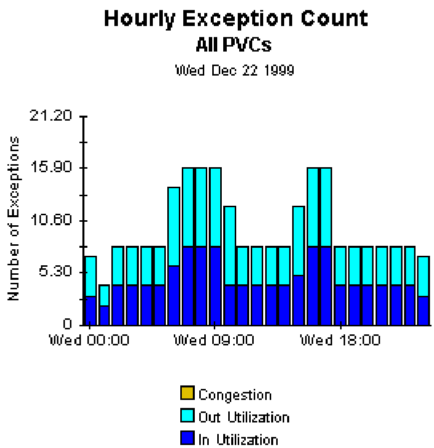

Forecast Report

The Forecast report lists all PVCs that are within 90 days of reaching a utilization or congestion threshold. Those PVCs closest to the 90-day threshold are listed first. The forecast data is generated using the daily 95th percentile values for in and out utilization from the daily exception table. The forecast is based on the data from the previous 42 days. The report uses only the 90-day forecast for a constraint. However, the table will display 30, 60, and 90-day forecasts. Drill-down charts display Grade of Service, Utilization, and Congestion information for that particular circuit over the baseline period (by default, 6 weeks).

Note: **For an in-depth analysis of how to interpret the Frame Relay Forecast report and for chart illustrations, see** “Forecast Report” on page 3-7.

The Forecast report comprises the following charts:

- The **Estimated Days to Threshold** table ranks the PVCs in ascending order that will reach a 80% In Utilization and Out Utilization threshold within 90 days.

  Double-click on a device to update the content of the following drill-down components with performance data for the device:

  - The **Grade of Service** stacked bar chart shows Grade of Service (GOS) scores for In Utilization, Out Utilization, and Congestion for the selected PVC for each day of the baseline period. Table 4-3 describes how Grade of Service is calculated.
  - The **Utilization** line graph shows In Utilization and Out Utilization, computed using both CIR and ifSpeed, for the selected PVC for each sample period on the report day.
  - The **Congestion** line graph shows congestion for the selected PVC for each sample taken on the report day.

QuickView and Snapshot Reports

The QuickView report identifies the top ten PVCs with the highest (worst) grade of service scores (according to Grade of Service score). Network managers and analysts can go to this report to select a PVC from a grade of service chart, and assess its status in several performance areas. Drill-down charts show daily distribution of bandwidth, as polled volume against the baseline and congestion at the as polled level.

Alternatively, the Snapshot report shows the same information, however, you select PVCs you want to view from a pick list when you invoke the report. For more information about selecting a PVC for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

Note: For an in-depth analysis of how to interpret the Frame Relay QuickView and Snapshot reports and for chart illustrations, see “QuickView and Snapshot Reports” on page 3-20**.**

The QuickView and Snapshot reports comprise the following charts:

- (QuickView) The **Daily Grade of Service** stacked bar chart shows the PVCs with the highest daily GOS scores. Table 4-3 describes how Grade of Service is calculated.
- (Snapshot) The **Snapshot Selection** table lists the time period for which the PVC information was collected, the PVC name, the interface, the PVC’s Input Utilization (CIR and ifSpeed), Output Utilization (CIR and ifSpeed), and Volume.

  Double-click on a device in the table to update the content of the following drill-down components with performance data for the device:

  - (Snapshot) The **Grade of Service** stacked bar chart shows the selected PVC’s GOS scores for each sample taken over the report day.
  - The **Hourly Volume** area graph shows volume for the selected PVC for each sample taken over the report day.
  - The **Hourly Congestion** line graph plots the number of frames with Forward Explicit Congestion Notification and the number of frames with Backward Explicit Congestion Notification for the selected PVC for each sample taken during the report day.
  - The **Distribution of Bandwidth Utilization - Input** pie chart shows the percentage of samples taken during the day in which the In Utilization metric calculated for the samples fall in a defined range. For more information about frequency distribution, see “Frequency Distribution” on page 2-14.
  - The **Distribution of Bandwidth Utilization - Output** pie chart shows the percentage of samples taken during the day in which the Out Utilization metric calculated for the samples fall in a defined range. For more

information about frequency distribution, see “Frequency Distribution” on page 2-14.

Near Real Time QuickView and Snapshot Reports

The Near Real Time--QuickView report provides PVC performance statistics up to the last SNMP poll. The report does not rely on nightly summaries; therefore, it provides instant reporting on collected data. Network managers and analysts can go to this report to select a PVC from a table, and assess its status in the following performance areas: utilization (CIR and ifSpeed), volume, and FECNs and BECNs.

Alternatively, the Snapshot report shows the same information, however, you select PVCs you want to view from a pick list when you invoke the report. For more information about selecting a PVC for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

The QuickView and Snapshot reports comprise the following charts:

- The **PVC Selection List** table shows each PVCs and its performance statistics based on the data polled for the last 6 hours. Depending on the type of statistic, the value shown is either the average or total value of all the data collected over the 6-hour period.

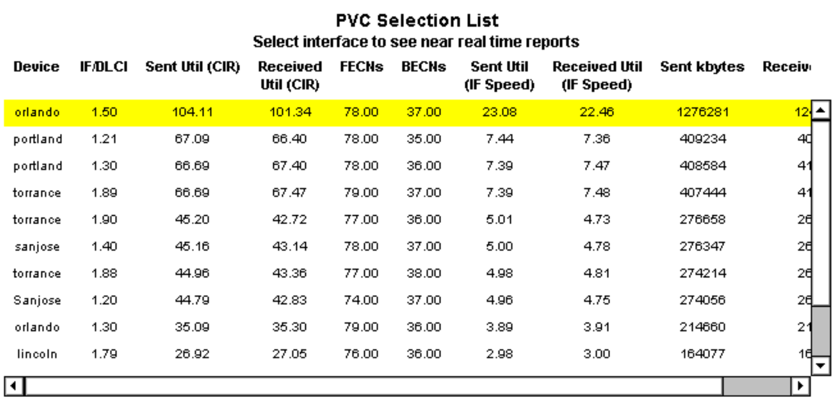
PVC Selection List
Select interface to see near real time reports

Device	IF/DLCI	Sent Util (CIR)	Received Util (CIR)	FECNs	BECNs	Sent Util (IF Speed)	Received Util (IF Speed)	Sent kbytes	Receiv
orlando	1.50	104.11	101.34	78.00	37.00	23.08	22.46	1276281	12
portland	1.21	67.09	66.40	78.00	35.00	7.44	7.36	409234	40
portland	1.30	66.69	67.40	78.00	36.00	7.39	7.47	408584	41
torrance	1.89	66.69	67.47	79.00	37.00	7.39	7.48	407444	41
torrance	1.90	45.20	42.72	77.00	36.00	5.01	4.73	276658	26
sanjose	1.40	45.16	43.14	78.00	37.00	5.00	4.78	276347	26
torrance	1.88	44.96	43.36	77.00	38.00	4.98	4.81	274214	26
Sanjose	1.20	44.79	42.83	74.00	37.00	4.96	4.75	274056	26
orlando	1.30	35.09	35.30	79.00	36.00	3.89	3.91	214660	21
lincoln	1.79	26.92	27.05	76.00	36.00	2.98	3.00	164077	16

Double-click on a device in the table to update the content of the following drill-down components with performance data for the device:

- The **PVC Utilization Based on CIR** line graph shows the selected PVC's utilization (CIR) for each sample taken during the report day.

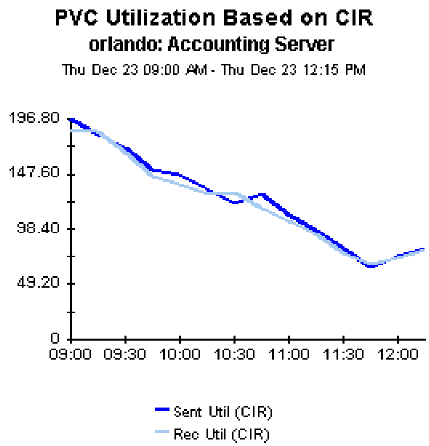

- The **PVC Volume** line graph shows volume for the selected PVC for each sample taken during the report day.

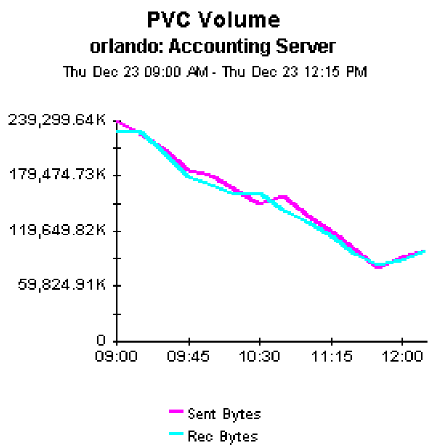

- The **PVC Utilization Based on Interface Speed** line graph shows the selected PVC's utilization (ifSpeed) for each sample taken during the report day.

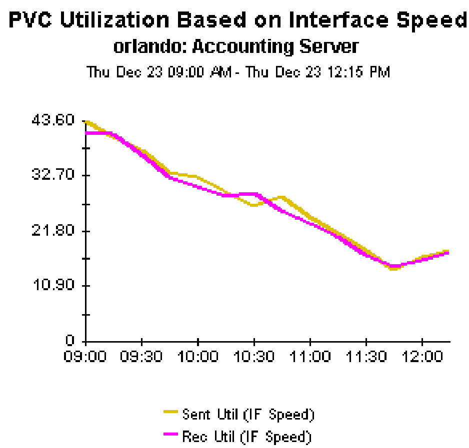

- The **PVC FECNs/BECNs** line graph plots the number of frames with Forward Explicit Congestion Notification and the number of frames with Backward Explicit Congestion Notification for the selected PVC for each sample taken during the report day.

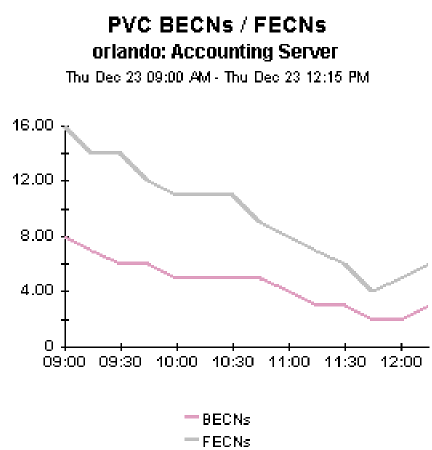

Capacity Planning Report

The Capacity Planning report shows the most overutilized and underutilized PVCs based on a 90-day forecasted value. Overutilized and underutilized PVCs are contrasted on this report to point out the feasibility of routing some traffic from the overutilized PVCs to the underutilized PVCs to balance the workload. Drill-down charts show daily data for average in and out utilization, grade of service, and congestion for the previous 42 days.

The Capacity Planning report comprises the following:

- The **Overutilized PVCs** table ranks the PVCs in descending order by In Utilization and Out Utilization (70% threshold) for 30, 60, and 90 days. Table 4-2 describes how In and Out Utilization are computed.

Overutilized PVCs

Projected To Exceed 70% Input or Output Utilization Within 90 Days

All Values Based On 95th Percentile

Device	Interface/DLCI	Description	Projected In Util 30 Days	Projected Out Util 30 Days	Projected In Util 60 Days	Projected Out Util 60 Days	Projectec Util 90 Da
orlando	1.50	Accounting Server	366.26	371.85	371.64	378.35	377.01
portland	1.21	Accounting	186.61	186.30	199.93	199.16	213.26
torrance	1.89	Support	184.79	187.06	197.10	200.43	209.42
portland	1.30	Sales	185.73	186.24	198.38	199.10	211.02
torrance	1.90	Marketing	96.54	103.49	110.99	116.84	125.44
Sanjose	1.20	Support	95.33	104.05	109.08	117.71	122.82
torrance	1.88	Sales	95.19	103.56	108.68	116.94	122.18
sanjose	1.40	Marketing	95.55	102.92	109.48	115.88	123.42
orlando	1.30	Accounting	65.36	64.72	79.63	78.75	93.89

- The **Underutilized PVCs** table ranks the PVCs projected to be the least utilized in terms of In Utilization or Out Utilization within the next 90 days in ascending order.

Underutilized PVCs
Projected To Be Least Utilized In 90 Days
All Values Based On 95th Percentile (P95)

Device	Interface/DLCI	Description	Current P95 In Utilization	Projected In Util 30 Days	Projected In Util 60 Days	Projected In Util 90 Days	Curr U
COPBKR_BCI01	202104.0.19	Sales	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.23	Sales	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.24	Sales	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.27	Sales	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.29	Sales	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.22	Support	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.26	Support	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.18	Marketing	0.00	0.00	0.00	0.00	
COPBKR_BCI01	202104.0.28	Accounting	0.00	0.00	0.00	0.00	

Double-click on a device in either table to update the content of the following drill-down components with performance data for the device:

- The **Grade of Service** bar chart shows Grade of Service (GOS) scores for the selected PVC. Table 4-3 describes how Grade of Service is calculated.

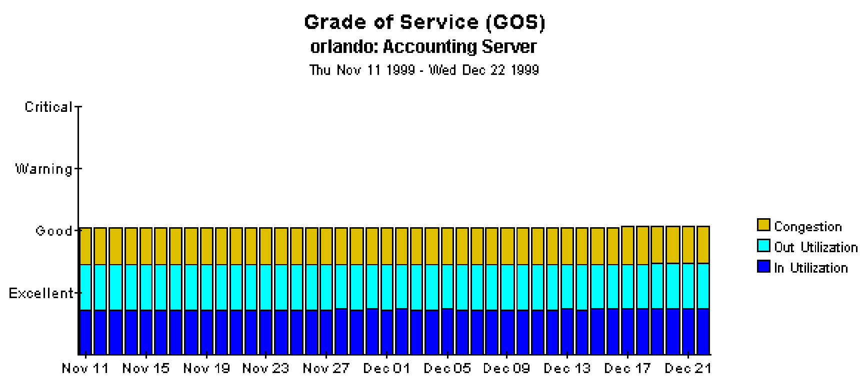

- The **Utilization** line graph shows In Utilization and Out Utilization calculated using CIR and ifSpeed for the selected PVC for each day of the baseline period.

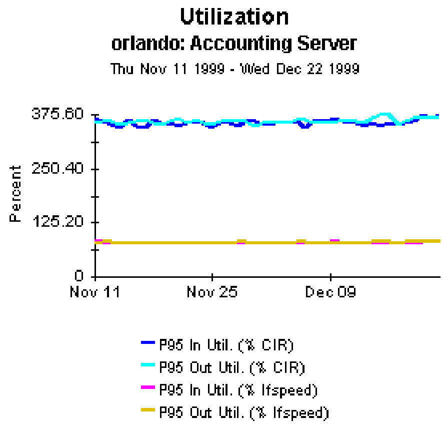

- The **Congestion** line graph shows congestion for the selected PVC for each day of the baseline period.

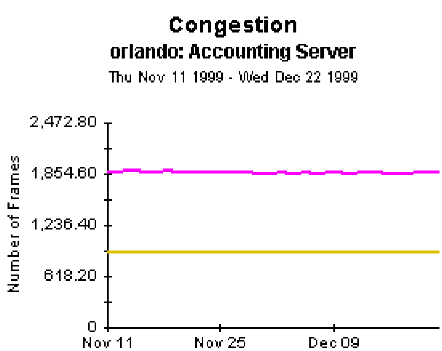

Hot Spots Report

The Hot Spots report provides the network manager and other members of the network management staff a list of PVCs experiencing utilization, congestion, or grade of service exceptions. The top ten offenders, ranked by total number of exceptions, are presented in a summary table to assist the network management team in analyzing the problems. Drill-down charts display grade of service, utilization, and congestion charts, as well as an exception detail table that shows the time of each exception and the value for the offending conditions as well as other categories.

The Hot Spots report comprises the following:

- The **Problem Summary for the Day** table ranks the PVCs with the highest exceptions counts ranked by total number of exceptions (input utilization + output utilization + congestion) in descending order.

Problem Summary For The Day
Frame Relay PVCs With Most Exceptions
Wed Dec 22 1999

Device	Element	Description	Congestion Exceptions	In Util Exceptions	Out Util Exceptions	Total Exceptions	Utilizatic Exceptio
orlando	1.50	Accounting Server	0.00	93.00	94.00	187.00	187.00
portland	1.30	Sales	0.00	81.00	81.00	162.00	162.00
torrance	1.89	Support	0.00	80.00	81.00	161.00	161.00
portland	1.21	Accounting	0.00	79.00	81.00	160.00	160.00
torrance	1.88	Sales	0.00	23.00	32.00	55.00	55.00
torrance	1.90	Marketing	0.00	20.00	35.00	55.00	55.00
Sanjose	1.20	Support	0.00	23.00	31.00	54.00	54.00
sanjose	1.40	Marketing	0.00	21.00	33.00	54.00	54.00

Double-click on a device to update the content of the following charts with performance data for the device:

- The **Grade of Service** stacked bar chart shows the Grade of Service (GOS) ratings for In Utilization, Out Utilization, and Congestion for the

selected PVC for each sample taken on the report day for the selected PVC. Table 4-3 describes how Grade of Service is calculated.

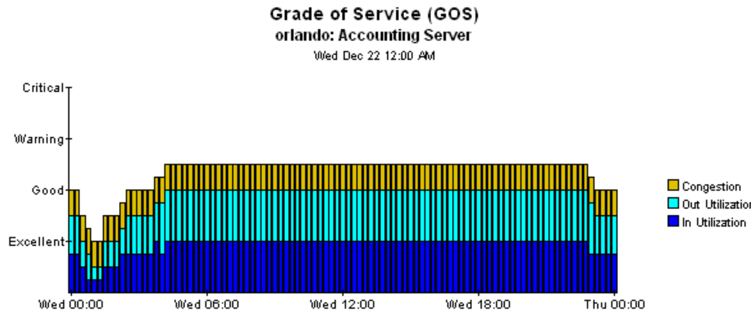

◆ The **Utilization** line graph shows In Utilization and Out Utilization, computed using both CIR and ifSpeed, for the selected PVC for each sample period on the report day.

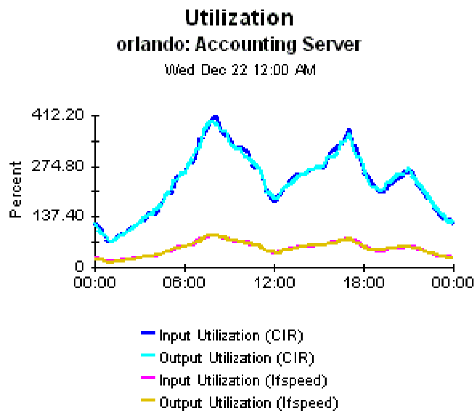

- The **Congestion** line graph shows congestion for the selected PVC for each sample taken on the report day.

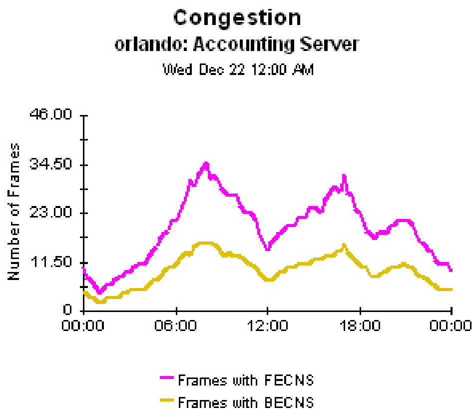

- The **Exception Detail** table shows the number of In Utilization, Out Utilization, and Congestion exception conditions occurring for each sample period on the report day for the selected PVC. An In Utilization or Out Utilization metric greater than 80% or Congestion metric greater than 25% is treated as an exception and is included in the chart.

Exception Detail
orlando: Accounting Server
Wed Dec 22 12:00 AM

Time Period	Interface/DLCI	Congestion (FECNs+BECNs)	In Util (CIR)	In Util (IfSpeed)	Out Util (CIR)	Out Util (IfSpeed)
Wed Dec 22 12:00 AM	1.50	0.00	121.24	26.87	120.53	26.71
Wed Dec 22 12:15 AM	1.50	0.00	112.56	24.95	104.36	23.13
Wed Dec 22 12:30 AM	1.50	0.00	95.76	21.23	91.79	20.35
Wed Dec 22 12:45 AM	1.50	0.00	79.00	17.51	87.25	19.34
Wed Dec 22 01:30 AM	1.50	0.00	80.40	17.82	80.95	17.94
Wed Dec 22 01:45 AM	1.50	0.00	89.96	19.94	86.51	19.18
Wed Dec 22 02:00 AM	1.50	0.00	94.84	21.02	95.21	21.10
Wed Dec 22 02:15 AM	1.50	0.00	100.42	22.26	99.92	22.15

Top Ten Report

The Top Ten report lets the network manager know at-a-glance which PVCs produce the greatest volume and the poorest health. The top ten elements, ranked by total volume and grade of service are listed by highest rank as well as highest change of rank. Volume and health change charts alert users to the PVC's greatest changes in volume, suggesting potential capacity overloading if increasing trends continue. The volume change report also may highlight which newly installed devices and technologies are approaching a logical threshold and which may be exceeding a threshold, or which are underutilized. Drastic increases in grade of service scores suggest additional PVCs to be watched, while decreases document an effective remedy to a previously negative condition.

The Top Ten Report comprises the following:

- The **PVCs with Highest Volume** table shows up to ten PVCs sorted by the current rank in ascending based on the report day's volume.

PVCs With Highest Volume

Wed Dec 22 1999

Device	Interface/DLCI	Description	Volume	Volume Rank	Previous Rank	Utilization	Congestion
orlando	1.50	Accounting Server	587821902.63	1	1	469.77	0.00
torrance	1.89	Support	156539485.82	2	3	250.22	0.02
portland	1.21	Accounting	156497759.61	3	4	249.83	0.02
portland	1.30	Sales	156121302.16	4	2	249.41	0.02
torrance	1.90	Marketing	87768136.22	5	8	140.13	0.04
Sanjose	1.20	Support	87762855.55	6	5	140.21	0.04
sanjose	1.40	Marketing	87698904.20	7	7	140.04	0.04
torrance	1.88	Sales	87585468.15	8	6	140.08	0.04
orlando	1.30	Accounting	60177092.68	9	9	96.14	0.08
lincoln	1.79	Support Server	35373290.12	10	11	56.52	0.77

- The **PVCs with Greatest Change in Volume** table shows the PVCs sorted by change in rank in ascending order based on the report day's volume change.

PVCs With Greatest Change In Volume
Wed Dec 22 1999

Device	Interface/DLCI	Description	Volume	Previous Volume	Volume Change Rank	Volume Rank	Pre
portland	1.30	Sales	156121302.16	156689312.71	1	4	
orlando	1.50	Accounting Server	587821902.63	588363836.02	2	1	
portland	1.21	Accounting	156497759.61	156096161.74	3	3	
torrance	1.90	Marketing	87768136.22	87419287.99	4	5	
torrance	1.89	Support	156539485.82	156390717.66	5	2	
torrance	1.88	Sales	87585468.15	87684576.72	6	8	
sanjose	1.40	Marketing	87698904.20	87625234.06	7	7	
orlando	1.30	Accounting	60177092.68	60137201.84	8	9	
Sanjose	1.20	Support	87762855.55	87744522.20	9	6	
atlanta	1.24	Marketing	35358425.49	35371012.58	10	11	

An advantage to using this table is that it may highlight which newly installed devices and technologies are approaching a logical threshold, may be exceeding a threshold, or may be impeding optimum utilization.

- The **PVCs with Highest Grade of Service (GOS) Scores** table shows PVCs sorted by rank in ascending order based on the report day's GOS Score. Table 4-3 describes how Grade of Service is calculated.

PVCs With Highest Grade Of Service (GOS) Scores
Wed Dec 22 1999

Device	Interface/DLCI	Description	GOS Score	GOS Rank	Previous Rank	Utilization	Congestion	V
lincoln	1.19	Accounting	2.16	1	1	52.77	4.19	330
atlanta	1.42	Marketing	2.15	2	1	52.75	4.19	330
orlando	1.50	Accounting Server	1.81	3	2	469.77	0.00	5878
portland	1.30	Sales	1.56	4	5	249.41	0.02	156
portland	1.21	Accounting	1.56	4	4	249.83	0.02	156
torrance	1.89	Support	1.55	5	3	250.22	0.02	156
torrance	1.90	Marketing	1.10	6	9	140.13	0.04	877
torrance	1.88	Sales	1.10	7	8	140.08	0.04	875
Sanjose	1.20	Support	1.09	8	6	140.21	0.04	877
sanjose	1.40	Marketing	1.09	9	7	140.04	0.04	876

- The **PVCs with Greatest Change in Grade of Service (GOS) Score** table shows the PVCs sorted by change in rank in ascending order based on the report day's GOS Score change. Drastic increases in grade of service scores

suggest additional PVCs be watched, while decreases document an effective remedy to a previously negative condition. Table 4-3 describes how Grade of Service is calculated.

PVCs With Greatest Change In Grade Of Service (GOS) Score
Wed Dec 22 1999

Device	Interface/DLCI	Description	GOS Score	Previous GOS Score	GOS Change Rank	GOS Rank	Previous Ranl
lincoln	1.19	Accounting	2.16	2.15	1	1	1
portland	1.30	Sales	1.56	1.55	2	4	5
torrance	1.89	Support	1.55	1.56	3	5	3
torrance	1.90	Marketing	1.10	1.09	4	6	9
Sanjose	1.20	Support	1.09	1.10	5	8	6
sanjose	1.40	Marketing	1.09	1.10	6	9	7
torrance	1.88	Sales	1.10	1.09	7	7	8
atlanta	1.24	Marketing	1.00	1.00	8	10	10
atlanta	1.42	Marketing	2.15	2.15	9	2	1
lincoln	1.79	Support Server	1.00	1.00	10	10	10

Service Level Management Report

The Service Level Management report let CIOs, CFOs, network managers, users and customers know how their Frame Relay service is performing against contracted service levels. It reports the ten PVCs with lowest availability and the ten PVCs with the highest network response times. It also provides charts for daily availability and daily and hourly response times for the elements with lowest availability or highest network response time.

The Service Level Management Report comprises the following:

- The **PVCs with Lowest Availability for the Day** bar chart shows PVCs with the lowest availability for the previous day. For more information about how availability is calculated, see "Availability Table" on page 3-42.

Availability
PVCs With Lowest Availability
Wed Dec 22 1999

Device	Interface/DCLI	Availability
Sanjose	1.20	100.00
portland	1.30	100.00
torrance	1.88	100.00
atlanta	1.24	100.00
atlanta	1.42	100.00
sanjose	1.40	100.00
torrance	1.89	100.00
lincoln	1.19	100.00
orlando	1.30	100.00
torrance	1.90	100.00
portland	1.21	100.00

Double-click on a device to update the content of the following drill-down components with performance data for the device:

- The **Daily Availability** line graph compares the availability of the selected PVC with the Service Level Agreement over the baseline period. By default, the service level is set at 99.5%. For more information about how

daily availability is calculated, see “Daily Availability Drill-Down” on page 3-42.

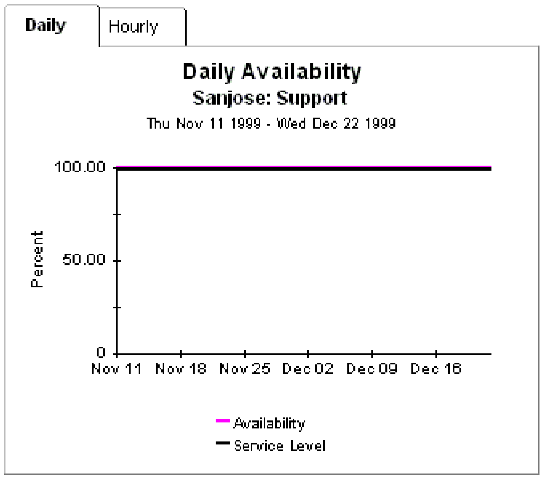

- The **Hourly Availability** line graph compares the availability of the selected PVC with the Service Level Agreement over the baseline period.

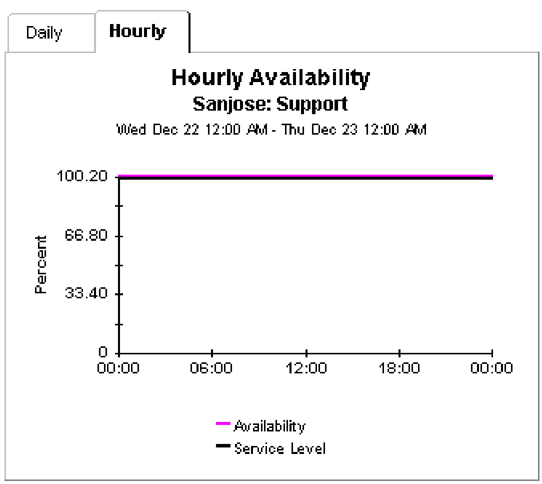

- The **PVCs with Highest Response Time for the Day** bar chart shows the PVCs with the highest (average) response times for the previous day. For more information about how network response times are calculated, see "Response Time" on page 3-45.

Response Time
PVCs With Highest Response Time
Wed Dec 22 1999

Device	Interface/DCLI	Resp. Time
COPBKR_BCI01	202104.0.27	219.00
COPBKR_BCI01	202104.0.23	216.00
COPBKR_BCI01	202104.0.17	216.00
COPBKR_BCI01	202104.0.24	215.00
COPBKR_BCI01	202104.0.16	215.00
COPBKR_BCI01	202104.0.19	214.00
COPBKR_BCI01	202104.0.25	214.00
COPBKR_BCI01	202104.0.26	213.00
COPBKR_BCI01	202104.0.22	212.00
COPBKR_BCI01	202104.0.28	212.00

Double-click on a device to update the content of the following drill-down components with performance data for the device:

- The **Daily Response Time** line graph shows the average response time for the selected PVC for each day of the baseline period. For more information about how daily response time is calculated, see "Daily Response Time Drill-Down" on page 3-46.

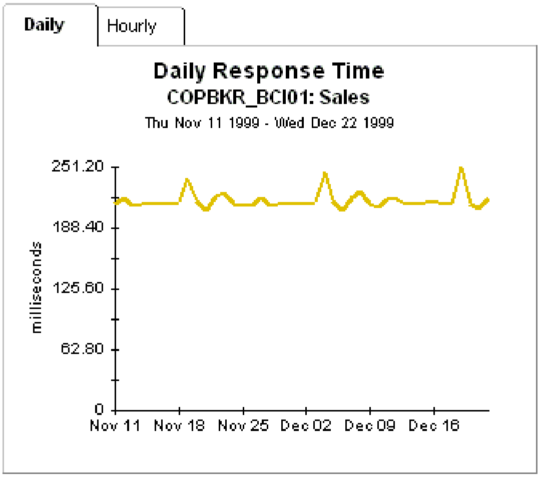

- The **Hourly Response Time** line graph shows the response time for each sample for the report day for the selected PVC. For more information about hourly network response time, see "Hourly Response Time Drill-Down" on page 3-47.

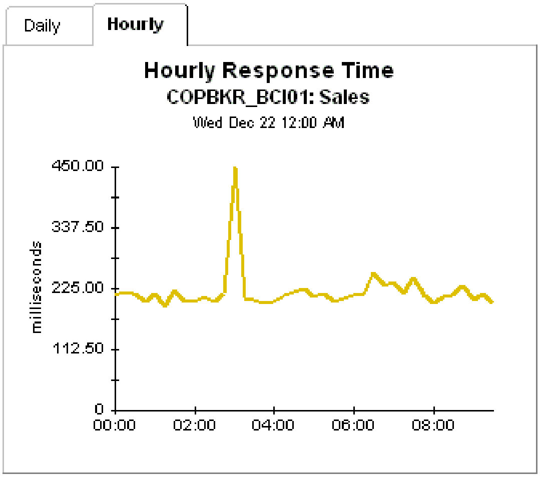

Data Source

Source data for the Frame Relay ReportPack is collected from the Frame Relay MIB: RFC 1315, which is supported by most routers, CSU/DSUs, and other network devices.

Table 4-1 lists the variables that are polled by SNMP from the RFC 1315 MIB.

Table 4-1: Frame Relay Variables

RFC 1315 Variable	OID String
frCircuitState	1.3.6.1.2.1.10.32.2.1.3
Indicates whether the particular virtual circuit is operational. In the absence of a Data Link Connection Management Interface, virtual circuit entries (rows) may be created by setting virtual circuit state to 'active', or deleted by changing Circuit state to 'invalid'. Whether or not the row actually disappears is left to the implementation, so this object may actually read as 'invalid' for some arbitrary length of time. It is also legal to set the state of a virtual circuit to 'inactive' to temporarily disable a given circuit.	
frCircuitReceivedFECNS	1.3.6.1.2.1.10.32.2.1.4
Number of frames received from the network indicating forward congestion since the virtual circuit was created.	
frCircuitReceivedBECNS	1.3.6.1.2.1.10.32.2.1.5
Number of frames received from the network indicating backward congestion since the virtual circuit was created.	
frCircuitSentFrames	1.3.6.1.2.1.10.32.2.1.6
The number of frames sent from this virtual circuit since it was created.	
frCircuitSentOctets	1.3.6.1.2.1.10.32.2.1.7
The number of octets sent from this virtual circuit since it was created.	

(1 of 2)

Table 4-1: Frame Relay Variables

RFC 1315 Variable	OID String
frCircuitReceivedFrames	1.3.6.1.2.1.10.32.2.1.8
Number of frames received over this virtual circuit since it was created.	
frCircuitReceivedOctets	1.3.6.1.2.1.10.32.2.1.9
Number of octets received over this virtual circuit since it was created.	
frCircuitLastTimeChange	1.3.6.1.2.1.10.32.2.1.11
The value of sysUpTime when last there was a change in the virtual circuit state.	

(2 of 2)

Basic Metric Calculations

The performance metrics used in Frame Relay reports are based on the calculations listed in Table 4-2. Note that some charts modify these calculations.

Table 4-2: Frame Relay Metrics

Metric	Computed As ...
In Utilization	(((inoctets*8)/delta_time)/CIR)*100 or (((inoctets*8)/delta_time)/ifSpeed)*100

(1 of 2)

Table 4-2: Frame Relay Metrics

Metric	Computed As ...
inoctets is multiplied by 8 to give the number of incoming bits, divided by delta_time to give bits-per-second, divided by Committed Information Rate (CIR) (which is collected as bits/sec), times 100 to yield a percentage. Interface speed (ifSpeed) is used if CIR is not available.	
Out Utilization	(((outoctets*8)/delta_time)/CIR)*100 or (((outoctets*8)/delta_time)/ifSpeed)*100
outoctets is multiplied by 8 to give the number of outgoing bits, divided by delta_time to give bits-per-second, divided by Committed Information Rate (CIR) (which is collected as bits/sec), times 100 to yield a percentage. Interface speed (ifSpeed) is used if CIR is not available.	
Congestion	((fecns+becns)/ (outframes+inframes))*100
Frames with Forward Explicit Congestion Notification (fecn) plus Frames with Backward Explicit Congestion Notification (becn), divided by the number of outframes plus inframes, times 100 to yield a percentage. fecns, becns, inframes, and outframes are collected as a number of frames since the virtual circuit was created. The number of frames with FECNs and BECNs, respectively, is collected from the frCircuitReceivedFECNs and frCircuitReceivedBECNs objects, respectively, in the Frame Relay MIB.	
Volume	total(inoctets + outoctets)
inoctets plus outoctets for the selected PVCs are totaled for the specified aggregation period.	

(2 of 2)

Grade of Service Calculation

Table 4-3 lists how the metrics are weighted in Frame Relay Grade of Service (GOS) charts.

Table 4-3: Frame Relay GOS Calculations

GOS Calculation	Input Utilization (CIR)	Output Utilization (CIR)	Congestion	Input Utilization (ifSpeed)	Output Utilization (ifSpeed)
Weighting	20%	20%	60%	20%	20%
"Excellent" Score Range	0-80%	0-80%	.50	0-40%	0-40%
"Good" Score Range	80-100%	80-100%	.50-1	40-60%	40-60%
"Warning" Score Range	100-150%	100-150%	1-2	60-80%	60-80%
"Critical" Score Range	Over 150%	Over 150%	Over 2	80-100%	80-100%

For more information about the Grade of Service metric, see Table 2-2 on page 2-6. For more information about Grade of Service as it applies to each type of report, see "Analyzing TREND Reports" on page 3-1.

TREND

5 LAN and WAN ReportPacks

The LAN (Local Area Network) and WAN (Wide Area Network) Connectivity ReportPacks direct polling of data from routers and other devices that support MIB-II to generate comprehensive report sets for the following:

- Half-duplex interfaces such as Ethernet, FDDI, and Token Ring interfaces (LAN)
- Full-duplex interfaces such as T-3 and DSL (WAN)

The LAN and WAN Connectivity ReportPacks answer questions including:

- Do the LAN or WAN networks provide an acceptable level of service to customers or end users?
- Are there overloaded LAN or WAN interfaces? If so, are there underutilized LAN or WAN interfaces that can help balance the load?
- Are performance problems due to high levels of congestion, errors, excessive utilization, or other causes?
- What are the top ten LAN or WAN interfaces contributing to the poor health of the network?
- Which LAN or WAN interfaces are likely to degrade service in the near future if preventive action is not taken?

Report Descriptions

The following sections describe each report in the LAN and WAN ReportPacks in detail.

Executive Summary Report

The Executive Summary report provides CFOs, CIOs, and other managers an overview of the performance of corporate interface in the critical areas of volume and utilization. Each chart shows key metrics aggregated for all interfaces. Key indicators of performance are shown individually and combined into a Grade of Service chart to reveal interface health at a glance.

Note: The TRENDweb Administrator can configure a LAN or WAN ReportPack installation, so that the Executive Summary reports will only display data for devices belonging to a specific customer. In Figure 3-1**, the report displays information for only the devices belonging to the Acme company. For more information about this feature, see** “Displaying Data for Specific Customer Devices in Reports” on page 1-20**.**

The Executive Summary report comprises the following components:

- The **Daily Volume** bar chart shows the total volume for all interfaces over the baseline period, giving the executive an instant view of increases or decreases in usage over time. By default, the baseline period consists of the six weeks prior to and including the report date.
- The **Hourly Volume** bar chart shows the total volume for each sample taken over the report period.
- The **Hourly Baseline** bar chart shows the average volume for each hour of the day over the report day. You can compare total volume for each hour of the report day with the baseline for that hour for the report day.

- The **Hourly Bandwidth Utilization** line chart shows the percent of bandwidth utilization for each sample taken over the report period.
- The **Hourly Grade of Service** stacked bar chart shows GOS scores for the LAN metrics (Utilization, Errors, and Discards) or WAN metrics (Input Utilization, Output Utilization, Total Percent Errors, and Total Percent Discards) over the baseline period for the selected interface. For more information about the GOS weighting calculations for LAN and WAN interfaces, see Table 5-4 and Table 5-5 respectively. Calculated values that are less than or equal to 1 show excellent performance.

Note: For an in-depth analysis of how to interpret the LAN or WAN Executive Summary report and for chart illustrations, see "Executive Summary Report" on page 3-2**.**

Forecast Report

The Forecast report lists interfaces that are within 90 days of reaching a metric threshold condition, indicating where near-term correction can prevent service degradation. For LAN interfaces, the metric thresholds are as follows: utilization (40%), discard (5%), or error (5%). For WAN interfaces, the metric thresholds are input and output utilization (80%). The calculation for the threshold is from the baseline period. Those interfaces closest to the 90-day threshold are listed first. The forecast data is generated using the daily 95th percentile values for utilization from the daily exception table.

The Forecast report comprises the following components:

◆ The **Estimated Days to Threshold (DTT)** summary table ranks interfaces in descending order by 90-day projected utilization.

Estimated Days To Threshold (DTT)
LAN Interfaces Exceeding or Within 90 Days Of Thresholds
All Values Based On 95th Percentile --- Thresholds:Util = 40%; Errors = 5%; Discards = 5%

Device	Interface	Description	DTT Utilization	Current P95 Utilization	Projected Util 90 Days	DTT Percent Errors	Current P95 Percent Errors
sanjose	1	"Sales"	51.00	34.63	44.15	-90.00	6.06
sanluis	2	"Corporate Offices"	52.00	34.37	43.83	-104.00	6.07
miami	1	"Sales"	53.00	34.63	43.82	-114.00	6.03
miami	2	"Mktgs"	54.00	34.41	43.57	-92.00	6.09
sanjose	2	"Mktgs"	105.00	30.88	38.76	-88.00	6.12

Double-click on an interface to update the content of the following drill-down components with performance data for the selected interface:

◆ The **Daily Grade of Service** stacked bar chart shows GOS scores for the LAN metrics (Utilization, Errors, and Discards) and WAN metrics (Input Utilization, Output Utilization, Total Percent Errors, and Total Percent Discards) over the baseline period for the selected interface. For more information about the GOS weighting calculations for LAN and WAN interfaces, see Table 5-4 and Table 5-5 respectively.

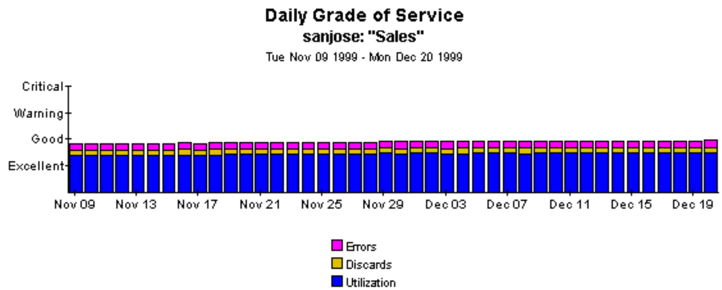

- The **Daily Utilization** line graph shows the average and 95th percentile utilization (LAN) or input and output utilization (WAN) for the selected interface for each day of the baseline period.

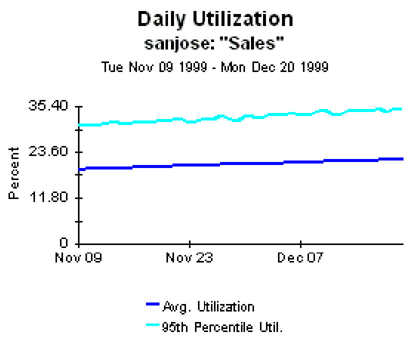

- The **Daily Errors and Discards** line graph shows the average percent and 95th percentile errors and discards for the selected interface for each day of the baseline period.

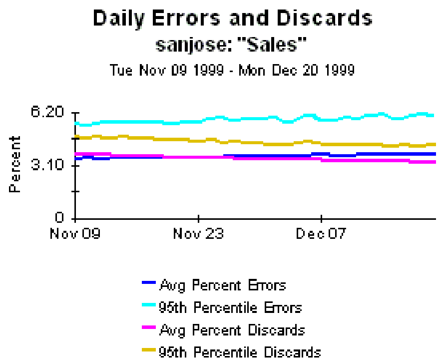

QuickView and Snapshot Reports

The QuickView report provides the network management staff with detailed information on key metrics for individual interfaces with the highest GOS scores.

The Snapshot report provides the same information as the QuickView report for individual interfaces that you select from a pick list for inclusion in the report. For more information about selecting a interface for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

QuickView

The QuickView report comprises the following components:

- The **Daily Grade of Service** stacked bar chart shows the GOS scores for the selected interfaces. The chart shows weighted values for the metrics for the report day. For LAN interfaces, the metrics are as follows: Utilization, Errors, and Discards. For WAN interfaces, the metrics are as follows: Input Utilization, Output Utilization, Total Percent Errors, and Total Percent Discards. For a description of how GOS is calculated for LAN and WAN interfaces, see Table 5-4 and Table 5-5 respectively.

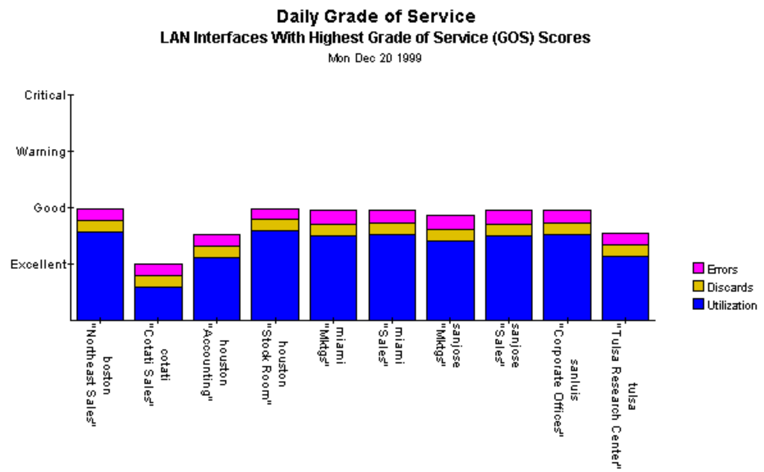

Double click on an interface in the chart update the following components with performance data for the selected interface:

- (LAN only) The **Hourly Volume** bar chart shows the volume of data flow for each sample period of the report day for the selected LAN interface. An overlay line graph shows the volume computed for each hour of the over the baseline period, so you can compare the volume shown for the report day with the baseline.

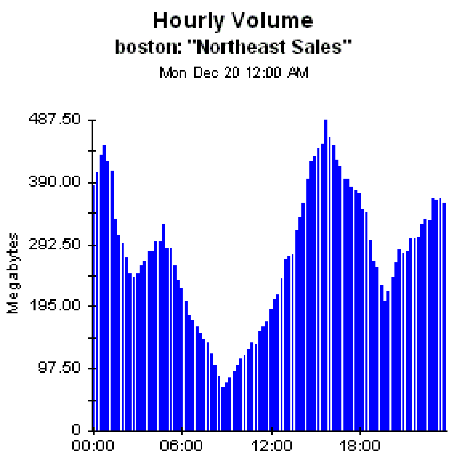

- (LAN only) The **Distribution of Bandwidth Utilization** pie chart shows the percentage utilization values that fall within the defined utilization buckets for the selected LAN interface for the report day.

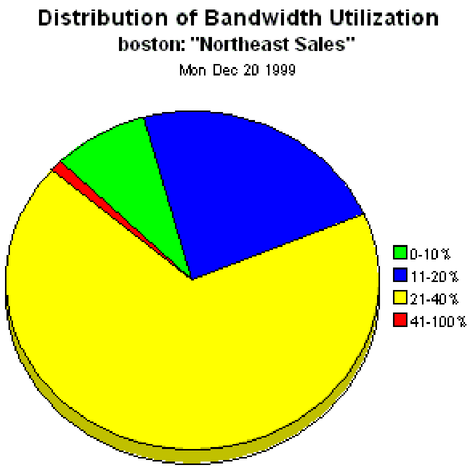

- (WAN only) The **Distribution of Input Bandwidth Utilization** pie chart shows show the percentage input utilization values that fall within the defined utilization buckets for the selected WAN interface for the report day.

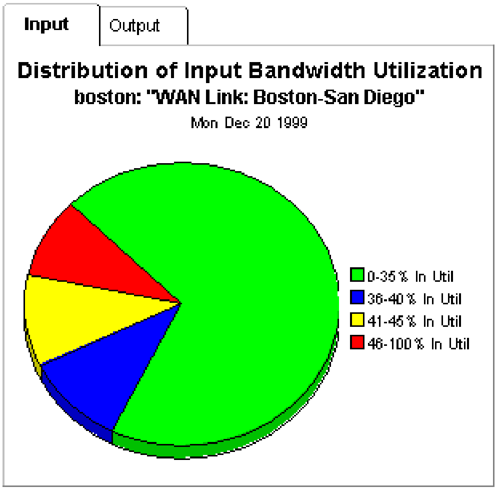

- (WAN only) The **Distribution of Output Bandwidth Utilization** pie chart shows the percentage output utilization values respectively that fall within the defined utilization buckets for the selected WAN interface for the report day.

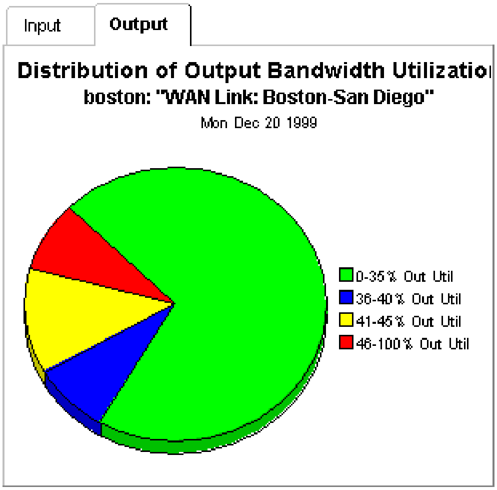

- The **Percent Errors and Percent Discards** line graph shows the percent errors and percent discards for the selected interface over the samples for each hour of the report day.

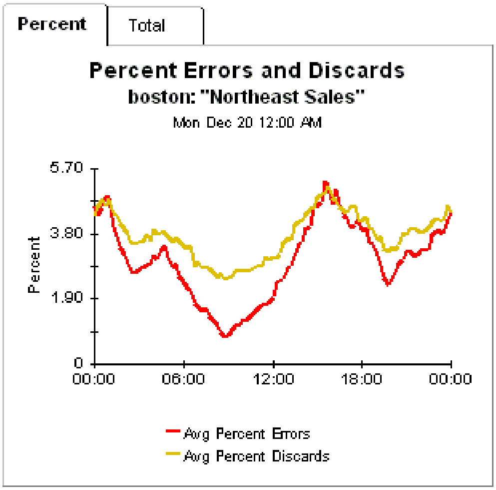

- (LAN only) The **Total Errors and Discards** line graph shows the total errors and total discards for the selected interface over each hour of the report day.

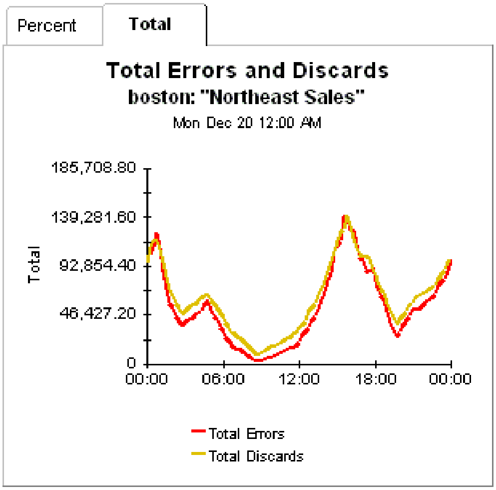

- The **Availability** line graph shows the percent availability of the selected interface over each hour of the report day. Availability is essentially computed as the ratio of the time that an interface is functional to the total time the interface is up, displayed as a percentage. See “Daily Availability Drill-Down” on page 3-42 for a discussion of the factors that are used to compute availability.

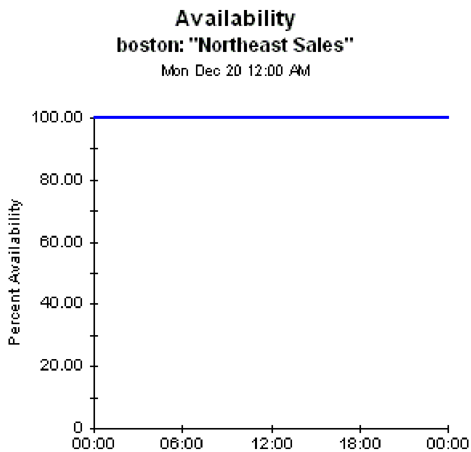

Snapshot

The Snapshot report comprises the following components:

- The **Snapshot Selection** table lists the selected interface and its metric information.

Snapshot Selection

Mon Dec 20 1999 - Mon Dec 20 1999

Time Period	Device	Element	Description	Utilization	Discards	Errors	Bytes
Mon Dec 20 12:00 AM	miami	1	"Sales"	34.30	7170648.82	8714305.00	229526321

Double click on an interface in the table to update the following components with performance data for the selected interface:

- The **Grade of Service** stacked bar chart shows the GOS scores for the selected interfaces. The chart shows weighted values for the metrics for the report day. For LAN interfaces, the metrics are as follows: Utilization, Errors, and Discards. For WAN interfaces, the metrics are as follows: Input Utilization, Output Utilization, Percent Errors, and Percent Discards. For a description of how GOS is calculated for LAN and WAN interfaces, see Table 5-4 and Table 5-5 respectively.
- (LAN only) The **Hourly Volume** bar chart shows the volume of data flow for each sample period of the report day for the selected LAN interface. An overlay line graph shows the volume computed for each hour of the over the baseline period, so you can compare the volume shown for the report day with the baseline.
- (LAN only) The **Distribution of Bandwidth Utilization** pie chart shows the percentage utilization values that fall within the defined utilization buckets for the selected LAN interface for the report day. The following table lists the buckets for input utilization
- (WAN only) The **Distribution of Input Bandwidth Utilization** and **Distribution of Output Bandwidth Utilization** pie charts show the percentage input utilization values and output utilization values respectively that fall within the defined utilization buckets for the selected WAN interface for the report day. The following table lists the buckets for input and output utilization.
- (LAN only) The **Percent Errors and Percent Discards** line graph shows the percent errors and percent discards for the selected interface over the samples for each hour of the report day.
- The **Total Errors and Discards** line graph shows the total errors (ifouterrors + ifinerrors + inunknownprotos) and total discards (ifindiscards + ifoutdiscards) for the selected interface over each hour of the report day.
- The **Availability** line graph shows the percent availability of the selected interface over each hour of the report day. Availability is essentially computed as the ratio of the time that an interface is functional to the total time the interface is up, displayed as a percentage. See “Daily Availability Drill-Down” on page 3-42 for a discussion of the factors that are used to compute availability.

Near Real Time QuickView and Snapshot Reports

The Near Real Time--QuickView and Near Real Time--Snapshot reports provide performance statistics up to the last SNMP poll. The report does not rely on nightly summaries; therefore, it provides instant reporting on collected data for individual LAN and WAN interfaces.

The Snapshot report provides the same information as the QuickView report for individual interfaces that you select from a pick list for inclusion in the report. For more information about selecting a interface for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

Near Real Time--QuickView

The QuickView report comprises the following components:

- The **Interface Selection List** table lists the available interfaces. and its performance statistics based on the data polled for the last 6 hours. The statistic shown for each metric is the average of all the data collected for that statistic over the 6-hour period.

Interface Selection List
Select interface to see near real time reports

Device	Interface	Speed	Avg Utilization	Peak Utilization	Avg Pct Discards	Avg Pct Errors	Description
boston	2	10000000	12.00	12.15	2.82	1.49	"Northeast Sales"
houston	3	10000000	11.99	12.08	2.05	0.00	"Stock Room"
miami	2	9000000	10.65	11.10	2.58	1.84	"Mktgs"
sanjose	1	9000000	10.48	10.94	2.56	1.82	"Sales"
miami	1	9000000	10.29	10.76	2.59	1.91	"Sales"
sanluis	2	9000000	10.27	10.70	2.64	1.91	"Corporate Offices"
sanjose	2	10000000	9.33	9.72	2.53	1.85	"Mktgs"
tulsa	2	10000000	6.63	6.73	1.98	0.00	"Tulsa Research Cent
houston	2	10000000	6.44	6.46	2.01	0.00	"Accounting"
cotati	2	10000000	1.92	1.95	1.19	0.10	"Cotati Sales"

Double click on an interface in the chart update the following components with performance data for the selected interface:

- The **Interface Average Traffic Rate** line graph shows the average traffic rate for each sample taken during the report period for the selected LAN interface.

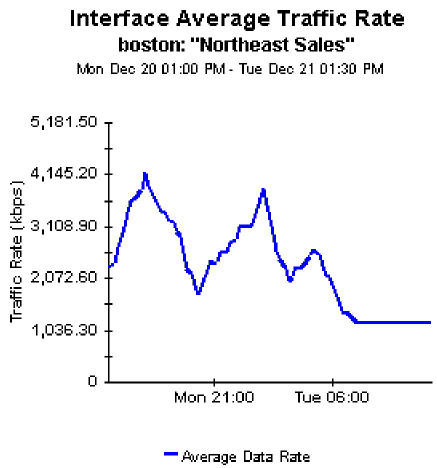

- The **Interface Utilization** line graph shows the average utilization for the percentage utilization for each sample taken during the report period.

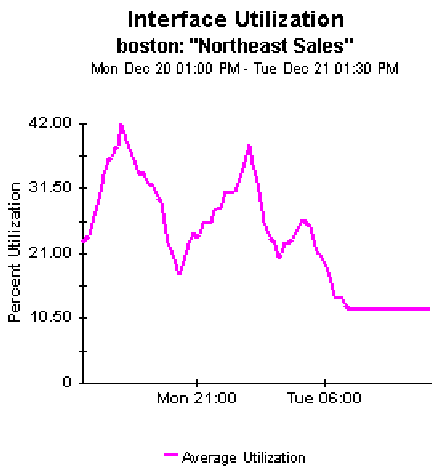

- The **Errors and Discards** line graph shows the total errors and total discards for the selected interface for each sample taken during the report period.

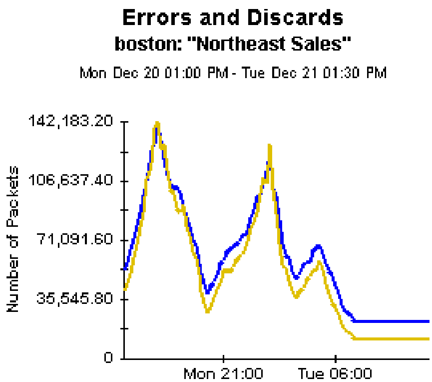

- The **Errors and Discards Percentages** line graph shows the percent errors and percent discards for the selected interface for each sample taken during the report period.

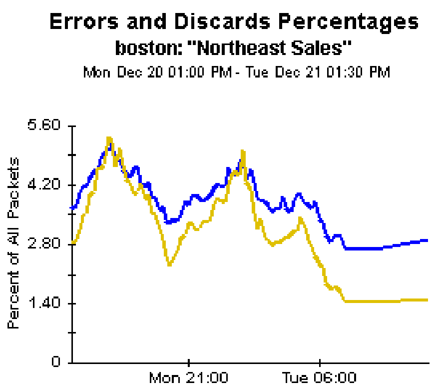

Near Real Time--Snapshot

The Snapshot report comprises the following components:

- The **Interface Selection List** table lists the available interfaces.

  Double click on an interface in the chart update the following components with performance data for the selected interface:

  - The **Interface Average Traffic Rate** line graph shows the average traffic rate for each sample period of the report day for the selected LAN interface.
  - The **Interface Utilization** line graph shows the average utilization for the percentage utilization for each sample period of the report day.
  - The **Errors and Discards** line graph shows the total errors and total discards for the selected interface for each sample period of the report day.
  - The **Errors and Percent Discards Percentages** line graph shows the percent errors and percent discards for the selected interface for each sample period of the report day.

Capacity Planning Report

The Capacity Planning report shows the most overutilized and underutilized interfaces based on a 90-day forecast value. It allows the CIO and network manager to estimate utilization levels in 30, 60, or 90 days.

Note: For an in-depth analysis of how to interpret the LAN or WAN Capacity Planning report and for chart illustrations, see "Capacity Planning Report" on page 3-13**.**

The Capacity Planning report comprises the following components:

- The **Overutilized Interfaces** summary table shows the interfaces projected to be 40% or higher utilized for LAN interfaces or 80% or higher input or output utilization (whichever is larger) for WAN interfaces within the next 90 days ranked in descending order by the 90-day projected utilization value. All calculations are based on 95th percentile utilization values and are displayed as

percentages in this table. For more information about how utilization is computed, see Table 5-2.

- The **Underutilized Interfaces** summary table shows the interfaces projected to be the least utilized within the next 90 days ranked in ascending order by the 90-day projected utilization value. All calculations are based on 95th percentile utilization values and are displayed as percentages in this table. Note that the associated drill-down components are the same for both underutilized and overutilized interfaces.

  Double-click on an interface in either table to update the following components with performance data for the selected interface:

  - The **Daily Grade of Service** stacked bar graph shows GOS scores for the LAN metrics (Utilization, Errors, and Discards) and WAN metrics (Input Utilization, Output Utilization, Total Percent Errors, and Total Percent Discards) over the baseline period for the selected interface. For more information about the GOS weighting calculations for LAN and WAN interfaces, see Table 5-4 and Table 5-5 respectively.
  - The **Daily Utilization** line graph shows the average and 95th percentile utilization over each day of the baseline period for the selected interface.
  - The **Daily Errors and Discards** line graph shows the average and 95th percentile errors and discards for the selected interface for each day of the baseline period. Values are shown as percentages.

Hot Spots Report

The Hot Spots report provides the network manager and other members of the network management staff a list of interfaces that are exceeding threshold conditions.

The Hot Spots report comprises the following components:

- The **Problem Summary for the Day** summary table displays the top ten interfaces in descending order by total number of exceptions. An *exception* occurs when an interface exceeds a metric threshold in any sample taken during the report day.

Problem Summary for the Day

LAN Interfaces With Highest Exception Counts

Mon Dec 20 1999

Device	Description	Utilization Exceptions	Error Exceptions	Discard Exceptions	Total Exceptions
sanjose	"Sales"	0.00	23.00	0.00	23.00
sanjose	"Mktgs"	0.00	23.00	0.00	23.00
sanluis	"Corporate Offices"	0.00	21.00	0.00	21.00
miami	"Sales"	0.00	19.00	0.00	19.00
miami	"Mktgs"	0.00	19.00	0.00	19.00
boston	"Northeast Sales"	1.00	3.00	1.00	5.00
houston	"Stock Room"	3.00	0.00	0.00	3.00

Following are the exceptions and thresholds for the LAN and WAN interfaces:

- **LAN**—Utilization (40%), Errors (5%), and Discards (5%)
- **WAN—**Input Utilization (80%), Output Utilization (80%), In Errors (5%), Out Errors (5%), In Discards (5%), and Out Discards (5%)

Double-click on any interface listed in the table to update the following charts and graphs with performance data for the selected interface:

- The **Grade of Service** stacked bar graph shows GOS scores for the LAN metrics (Utilization, Errors, and Discards) and WAN metrics (Input Utilization, Output Utilization, Total Percent Errors, and Total Percent Discards) over the baseline period for the selected interface. For more information about the GOS weighting calculations for LAN and WAN interfaces, see Table 5-4 and Table 5-5 respectively.

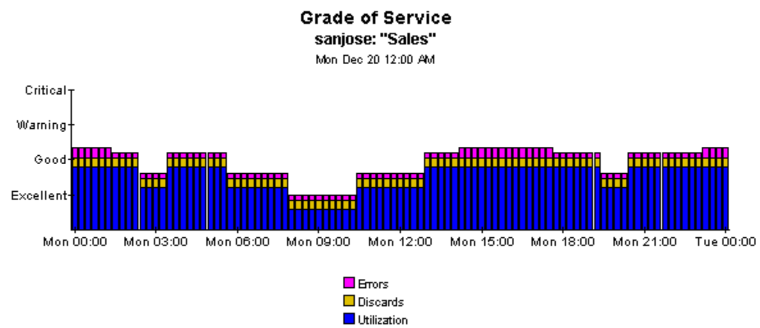

- The **Utilization** line graph shows the average percentage of bandwidth utilization (LAN) or the input and output utilization (WAN) for each hour of the report day for the specified interface.

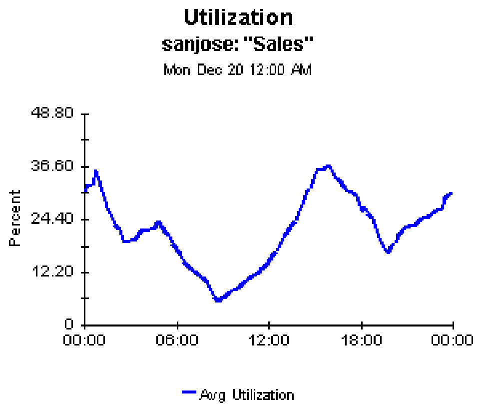

- The **Percent Errors and Discards** line graph shows the average percentage of errors and discards (LAN) or in errors, out errors, in

discards, and out discards (WAN) for the specified interface over the hours of the report day.

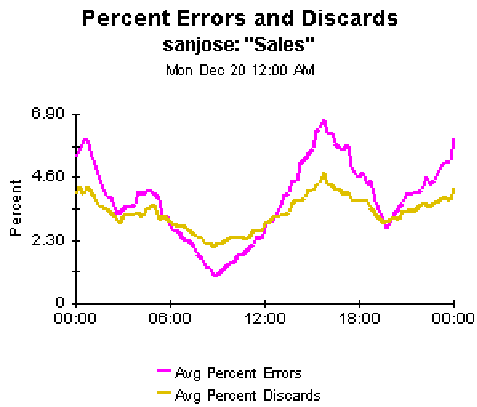

- The **Exception Detail** table shows the time the sample was taken for each exception and the values for the offending conditions. For more information about how the metrics are calculated for the LAN and WAN interfaces, see Table 5-2 and Table 5-3 respectively.

Exception Detail
sanjose: "Sales"
Mon Dec 20 12:00 AM

Time Period	Utilization	Percent Discards	Percent Errors	Pct Broadcast Packets	Avg Packet Size
Mon Dec 20 12:00 AM	30.57	4.00	5.29	0.10	112.84
Mon Dec 20 12:15 AM	32.66	4.28	5.54	0.11	121.60
Mon Dec 20 12:30 AM	32.31	3.98	5.94	0.09	107.29
Mon Dec 20 12:45 AM	35.92	4.27	6.04	0.10	116.84
Mon Dec 20 01:00 AM	32.34	4.06	5.38	0.10	113.57
Mon Dec 20 01:15 AM	30.05	3.86	5.11	0.10	115.98
Mon Dec 20 02:15 PM	28.94	3.76	5.08	0.10	116.08
Mon Dec 20 02:30 PM	31.07	3.87	5.24	0.10	115.22

Top Ten Report

The Top Ten report lets the network manager know at-a-glance which interfaces produce the greatest volume and the poorest health. The top ten elements, ranked by total volume and grade of service are listed by highest rank as well as highest change of rank. Volume and health change charts alert you to the interface's greatest changes in volume, suggesting potential capacity overloading if increasing trends continue. The volume change report also may highlight the newly installed devices and technologies that are approaching a utilization threshold or are underutilized. Drastic increases in GOS scores suggest additional interfaces to be watched, while decreases document an effective remedy to a previously negative condition. For more information about the GOS weighting calculations for LAN and WAN interfaces, see Table 5-4 and Table 5-5 respectively.

The Top Ten report comprises the following components:

- The **LAN Interfaces with Highest Volume** table shows up to ten interfaces with the highest volume for the report day.

LAN Interfaces With Highest Volume
Mon Dec 20 1999

Device	Interface	Description	Volume	Volume Rank	Previous Rank	Utilization
houston	3	"Stock Room"	290898951	1	1	23.79
boston	2	"Northeast Sales"	290036337	2	2	23.72
miami	2	"Mktgs"	240488729	3	6	21.84
sanjose	1	"Sales"	239835854	4	5	21.80
sanluis	2	"Corporate Offices"	239692493	5	4	21.79
miami	1	"Sales"	239041605	6	3	21.74
sanjose	2	"Mktgs"	239031279	7	7	19.57
houston	2	"Accounting"	167047995	12	13	13.65
tulsa	2	"Tulsa Research Center"	165739125	13	12	13.54
cotati	2	"Cotati Sales"	46660616	18	18	3.82

- The **LAN Interfaces with Greatest Change in Volume** table shows the current day's volume, previous day's volume, and rankings for up to ten interfaces with the greatest changes in volume. It ranks the interfaces by greatest change in volume. An advantage to using this table is that it may highlight the newly installed devices and technologies that are approaching a logical threshold, exceeding a threshold, or impeding optimum utilization.

LAN Interfaces With Greatest Change In Volume
Mon Dec 20 1999

Device	Interface	Description	Volume	Previous Volume	Volume Change Rank	Volume Rank	Previc
miami	2	"Mktgs"	240488729	238579410	1	3	
houston	3	"Stock Room"	290898951	289770889	4	1	
miami	1	"Sales"	239041605	240028759	5	6	
tulsa	2	"Tulsa Research Center"	165739125	166705624	6	13	
sanjose	1	"Sales"	239835854	238966889	8	4	
sanjose	2	"Mktgs"	239031279	238200269	9	7	
houston	2	"Accounting"	167047995	166292234	11	12	
boston	2	"Northeast Sales"	290036337	289608971	15	2	
cotati	2	"Cotati Sales"	46660616	46543182	17	18	
sanluis	2	"Corporate Offices"	239692493	239715997	20	5	

- The **LAN Interfaces with Highest Grade of Service (GOS) Scores** table ranks up to ten interfaces having the highest (and therefore poorest) GOS scores. Volume and utilization for each interface is also shown.

LAN Interfaces With Highest Grade Of Service (GOS) Scores
Mon Dec 20 1999

Device	Interface	Description	GOS Score	GOS Rank	Previous Rank	Utilization	Volume
houston	3	"Stock Room"	1.99	7	10	23.79	290898951
boston	2	"Northeast Sales"	1.98	8	9	23.72	290036337
sanluis	2	"Corporate Offices"	1.96	9	12	21.79	239692493
miami	1	"Sales"	1.96	10	11	21.74	239041605
sanjose	1	"Sales"	1.95	11	13	21.80	239835854
miami	2	"Mktgs"	1.95	12	14	21.84	240488729
sanjose	2	"Mktgs"	1.87	13	15	19.57	239031279
tulsa	2	"Tulsa Research Center"	1.54	16	19	13.54	165739125
houston	2	"Accounting"	1.52	17	18	13.65	167047995
cotati	2	"Cotati Sales"	1.00	18	20	3.82	46660616

- The **LAN Interfaces with Greatest Change in Grade of Service Scores** table shows the current and previous GOS scores for up to ten interfaces with the highest percentage of change. It ranks the interfaces by highest GOS percentage of change. Drastic increases in GOS scores changes identify interfaces that need to be monitored, while GOS score decreases document an effective remedy to a previously negative condition.

LAN Interfaces With Greatest Change In Grade Of Service (GOS) Scores
Mon Dec 20 1999

Device	Interface	Description	GOS Score	Previous GOS Score	GOS Change Rank	GOS Rank	Previous Rank	U
houston	2	"Accounting"	1.52	1.54	1	17	18	
houston	3	"Stock Room"	1.99	1.97	2	7	10	
tulsa	2	"Tulsa Research Center"	1.54	1.53	4	16	19	
miami	1	"Sales"	1.96	1.97	8	10	11	
sanjose	2	"Mktgs"	1.87	1.86	11	13	15	
sanjose	1	"Sales"	1.95	1.95	12	11	13	
miami	2	"Mktgs"	1.95	1.95	13	12	14	
boston	2	"Northeast Sales"	1.98	1.98	14	8	9	
sanluis	2	"Corporate Offices"	1.96	1.96	16	9	12	
cotati	2	"Cotati Sales"	1.00	1.00	19	18	20	

Service Level Management Report

The Service Level Management report lets CIOs, CFOs, network managers, and customers know if their interfaces are meeting contracted service levels for availability and response times by reporting interface availability across the enterprise over time and for individual interfaces for the day. It reports the ten interfaces with lowest availability for the day and the ten interfaces with the highest network response times. It also provides drill-down reports for daily availability over the baseline period and daily and hourly response times for the selected interfaces with lowest availability or highest network response time.

The Service Level Management report comprises the following components:

- The **LAN Interfaces with Lowest Availability for the Day** bar graph shows the interfaces with the lowest availability for the report day. Availability is essentially computed as the ratio of the time that an interface is functional to the total time the interface is up, displayed as a percentage. See "Daily Availability Drill-Down" on page 3-42 for a discussion of the factors that are used to compute availability.

Availability

LAN Interfaces With Lowest Availability For The Day

Mon Dec 20 1999

Device	Interface	Availability
miami	1	100.00
miami	2	100.00
sanjose	1	100.00
sanjose	2	100.00
cotati	2	100.00
houston	2	100.00
houston	3	100.00
boston	2	100.00
tulsa	2	100.00
sanluis	2	100.00

Double-click on a interface in this chart to update the following chart with information for the selected interface:

- The **Daily Availability** line graph shows the average daily availability for the selected interface for each day of the report period.

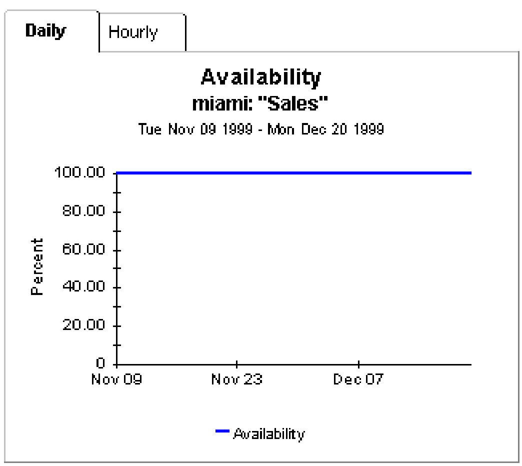

- The **Hourly Availability** line graph shows the average hourly availability of the selected interface for each hour for each sample taken during the report period.

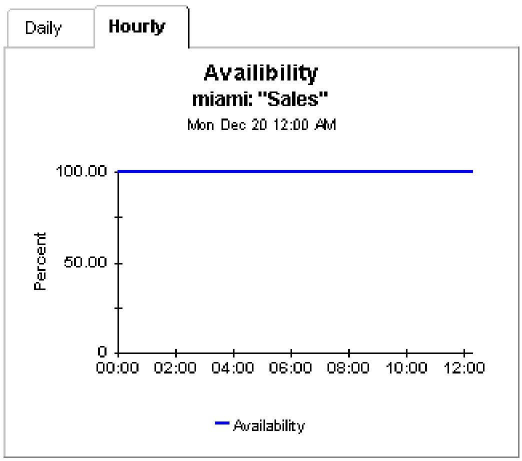

- The **LAN Interfaces with Highest Response Time for the Day** bar graph shows the interfaces with the highest (average) response times for the day in descending sequence from left to right. Double-click on a interface in this chart to update the following charts with detailed information for the selected interface:

Response Time

Highest SNMP Response Time by Interface

Mon Dec 20 1999

Device	Interface	Response Time
miami	1	98.00
miami	2	98.00
sanjose	1	98.00
sanjose	2	98.00
cotati	2	98.00
houston	2	98.00
houston	3	98.00
boston	2	98.00
tulsa	2	98.00
sanluis	2	98.00

- The **Daily Response Time** line graph shows the average daily network response time for the selected interface for each day of the baseline period. (Response time is taken from the database column received_usec, which is the time, in milliseconds, between the first SNMP GET and last GET RESPONSE request. If multiple packets are needed to get the entire MIB, the time for all packets is included.)

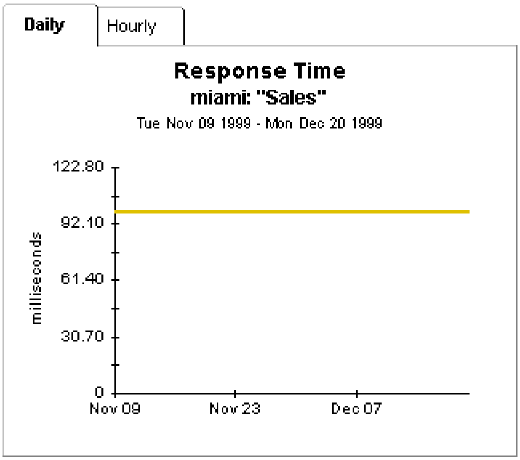

- The **Hourly Response Time** line graph shows the response time (received_usec) for each sample for the report day for the selected interface.

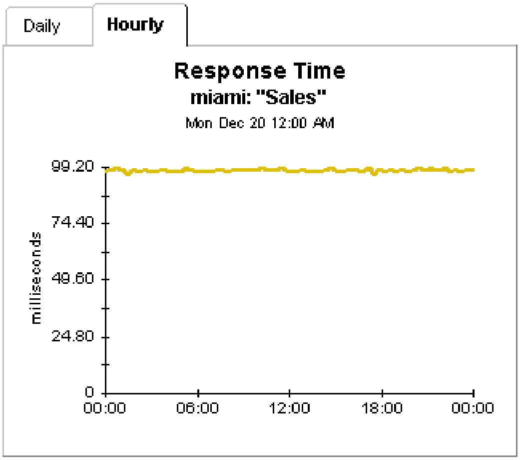

Data Source

Source data for the LAN and WAN ReportPacks is collected from the MIB II ifEntry table. Table 5-1 lists the variables

Table 5-1: LAN and WAN Connectivity Variables

Variable	OID String
ifIndex	1.3.6.1.2.1.2.2.1.1
A unique value for each interface. Its value ranges between 1 and the value of ifNumber. The value for each interface must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization.	

(1 of 5)

Table 5-1: LAN and WAN Connectivity Variables

Variable	OID String
ifDescr	1.3.6.1.2.1.2.2.1.2
A textual string containing information about the interface. This string should include the name of the manufacturer, the product name, and the version of the hardware interface.	
ifType	1.3.6.1.2.1.2.2.1.3
The type of interface, distinguished according to the physical/link protocol(s) immediately “below” the network layer in the protocol stack.	
ifMtu	1.3.6.1.2.1.2.2.1.4
The size of the largest datagram which can be sent/received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.	
ifSpeed	1.3.6.1.2.1.2.2.1.5
An estimate of the interface's current bandwidth in bits per second. For interfaces which do not vary in bandwidth or for those where no accurate estimation can be made, this object should contain the nominal bandwidth.	
ifPhysAddress	1.3.6.1.2.1.2.2.1.6
The interface's address at the protocol layer immediately “below” the network layer in the protocol stack. For interfaces which do not have such an address (e.g., a serial line), this object should contain an octet string of zero length.	
ifAdminStatus	1.3.6.1.2.1.2.2.1.7
The desired state of the interface. The testing (3) state indicates that no operational packets can be passed.	

(2 of 5)

Table 5-1: LAN and WAN Connectivity Variables

Variable	**OID String**
ifOperStatus	1.3.6.1.2.1.2.2.1.8
The current operational state of the interface. The testing (3) state indicates that no operational packets can be passed.	
ifLastChange	1.3.6.1.2.1.2.2.1.9
The value of sysUpTime at the time the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this object contains a zero value.	
ifInOctets	1.3.6.1.2.1.2.2.1.10
The total number of octets received on the interface, including framing characters.	
ifInUcastPkts	1.3.6.1.2.1.2.2.1.11
The number of subnetwork-unicast packets delivered to a higher-layer protocol.	
ifInNUcastPkts	1.3.6.1.2.1.2.2.1.12
The number of non-unicast (i.e., subnetwork-broadcast or subnetwork-multicast) packets delivered to a higher-layer protocol.	
ifInDiscards	1.3.6.1.2.1.2.2.1.13
The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.	
ifInErrors	1.3.6.1.2.1.2.2.1.14

(3 of 5)

Table 5-1: LAN and WAN Connectivity Variables

Variable	OID String
The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.	
ifInUnknownProtos	1.3.6.1.2.1.2.2.1.15
The number of packets received via the interface which were discarded because of an unknown or unsupported protocol.	
ifOutOctets	1.3.6.1.2.1.2.2.1.16
The total number of octets transmitted out of the interface, including framing characters.	
ifOutUcastPkts	1.3.6.1.2.1.2.2.1.17
The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent.	
ifOutNUcastPkts	1.3.6.1.2.1.2.2.1.18
The total number of packets that higher-level protocols requested be transmitted to a non-unicast (i.e., a subnetwork-broadcast or subnetwork-multicast) address, including those that were discarded or not sent.	
ifOutDiscards	1.3.6.1.2.1.2.2.1.19
The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.	
ifOutErrors	1.3.6.1.2.1.2.2.1.20
The number of outbound packets that could not be transmitted because of errors.	

(4 of 5)

Table 5-1: LAN and WAN Connectivity Variables

Variable	OID String
ifOutQLen	1.3.6.1.2.1.2.2.1.21
The length of the output packet queue (in packets).	
ifSpecific	1.3.6.1.2.1.2.2.1.22
A reference to MIB definitions specific to the particular media being used to realize the interface. For example, if the interface is realized by an ethernet, then the value of this object refers to a document defining objects specific to ethernet. If this information is not present, its value should be set to the OBJECT IDENTIFIER { 0 0 }, which is a syntactically valid object identifier, and any conformant implementation of ASN.1 and BER must be able to generate and recognize this value.	

(5 of 5)

Basic Metric Calculations

Table 5-2 lists the calculations for the performance metrics used in the LAN reports. Note that some charts modify these calculations.

Table 5-2: LAN Metrics

Metric	Computed As ...
Utilization	((ifoutoctets+ifinoctets)*8)/(delta_time/ifspeed+1)*100.0
This metric is the sum of inoctets and outoctets multiplied by 8 to give the number of bits, which is divided by total speed to give bits-per-second. Total bits-per-second is divided by ifspeed (which is collected as bits-per-second) plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.	

Table 5-2: LAN Metrics

Metric	**Computed As ...**
Errors	(ifouterrors+ifinerrors+inunknownprotos)/(ifoutucastpkts+ ifoutnucastpkts+inunknownprotos+ifinerrors+ifinucastpkts+ ifinnucastpkts+ifindiscards+1)*100.0
This metric is the ratio of the sum of inbound and outbound packets containing errors and those that are discarded because they come from an unknown source to the total packet traffic (inbound and outbound) at the interface plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.	
Discards	(ifindiscards+ifoutdiscards)/(ifinucastpkts+ifinnucastpkts+ ifindiscards+ifinerrors+inunknownprotos+ifoutnucastpkts+ ifoutucastpkts+1)*100.0
This metric is essentially the ratio of the sum of inbound and outbound packets that are discarded (for reasons other than containing errors) to the total packet traffic (inbound and outbound) at the interface plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.	
Volume	inoctets+outoctets
The inoctets and outoctets (bytes) for the selected managed elements are totaled for the specified aggregation period.	

Table 5-3 lists the calculations for the performance metrics used in WAN reports. Note that some charts modify these calculations.

Table 5-3: WAN Metrics

<table>
<tr><th>Metric</th><th>Computed As ...</th></tr>
<tr><td>Input Utilization</td><td>((ifinoctets*8)/delta_time/(ifspeed+1))*100.0</td></tr>
<tr><td colspan="2">inoctets are multiplied by 8 to give the number of bits, which is divided by delta_time to give bits-per-second. Bits-per-second is divided by ifspeed (which is collected as bits-per-second) plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.</td></tr>
<tr><td>Output Utilization</td><td>((ifoutoctets*8)/delta_time/(ifspeed+1))*100.0</td></tr>
<tr><td colspan="2">outoctets are multiplied by 8 to give the number of bits, which is divided by delta_time to give bits-per-second. Bits-per-second is divided by ifspeed (which is collected as bits-per-second) plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.</td></tr>
<tr><td>Utilization (combined Input and Output Utilization)</td><td>((ifoutoctets+ifinoctets)*8)/(delta_time)/(ifspeed+1)*100.0</td></tr>
<tr><td colspan="2">The sum of outoctets and inoctets is multiplied by 8 to give the number of bits, which is divided by delta_time to give bits-per-second. Bits-per-second is divided by ifspeed (which is collected as bits-per-second) plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.</td></tr>
<tr><td>Input Errors</td><td>(ifinerrors+inunknownprotos)/(ifinucastpkts+ ifinnucastpkts+ifindiscards+ifinerrors+ inunknownprotos)*100.0</td></tr>
</table>

Table 5-3: WAN Metrics

Metric	Computed As ...
The ratio of the number of inbound packets containing errors (ifinerrors) plus the number of inbound packets from an unsupported protocol (inunknownprotos) to all inbound packets regardless of type is multiplied by 100 to yield a percentage.	
Percent Output Errors	(ifouterrors)/(ifoutucastpkts+ifoutnucastpkts+1)*100.0
The ratio of the number of outbound packets containing errors (if outerrors) to the number of packets that higher-level protocols requested to be transmitted to a unicast or non-unicast address is multiplied by 100 to yield a percentage. (In this computation, a negligible amount, that is, 1 packet, is added to the divisor to prevent a divide-by-zero condition.)	
Percent Input Discards	(ifindiscards)/(ifinucastpkts+ifinnucastpkts+ ifindiscards+ifinerrors+inunknownprotos+1)*100.0
The ratio of the number of inbound packets that were discarded for reasons other than containing errors to all inbound packets regardless of type plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.	
Percent Output Discards	(ifoutdiscards)/(ifoutnucastpkts+ifoutucastpkts+1)*100.0
The ratio of the number of outbound packets that were discarded for reasons other than containing errors to the number of packets that higher-level protocols requested be transmitted to a unicast or non-unicast address plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.	
Volume	inoctets+outoctets
The inoctets and outoctets (bytes) for the selected WAN interfaces are totaled for the specified aggregation period.	
Total Errors	ifouterrors+ifinerrors+inunknownprotos

Table 5-3: WAN Metrics

Metric	**Computed As ...**
The sum of outbound packets that contain errors, inbound packets that contain errors, and inbound packets from unknown protocols.	
Total Discards	ifindiscards+ifoutdiscards
The sum of inbound packets that are discarded and outbound packets that are discarded.	
Percent Errors	(ifouterrors+ifinerrors+inunknownprotos)/ (ifoutucastpkts+ifoutnucastpkts+inunknownprotos+ ifinerrors+ifinucastpkts+ifinnucastpkts+ifindiscards+1) *100.0
The ratio of the sum of inbound and outbound packets containing errors and inbound packets from unknown protocols to the sum of inbound packets regardless of type and outbound ifoutucastpkts and ifoutnucastpkts plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.	
Percent Discards	(ifindiscards+ifoutdiscards)/ (ifinucastpkts+ifinnucastpkts+ifindiscards+ifinerrors+ inunknownprotos+ifoutnucastpkts+ifoutucastpkts+1)* 100.0
The ratio of the sum of inbound and outbound packets that are discarded to the sum of all inbound packets regardless of type and ifoutnucastpkts and ifoutucastpkts plus 1 packet (to prevent a potential divide by zero condition) multiplied by 100 to yield a percentage.	

Grade of Service Calculation

Table 5-4 lists how the metrics are weighted in LAN Grade of Service (GOS) charts.

Table 5-4: LAN GOS Calculations

GOS Calculation	Utilization	Errors	Discards
Weighting	60%	20%	20%
"Excellent" Score Range	0-9%	0-4%	0-4%
"Good" Score Range	10-19%	5-6%	5-6%
"Warning" Score Range	20-39	7-9%	7-9%
"Critical" Score Range	40% and over	10% and over	10% and over

Table 5-5 lists how the metrics are weighted in WAN Grade of Service (GOS) charts.

Table 5-5: WAN GOS Calculations

GOS Calculation	Input Utilization	Output Utilization	Total Percent Errors	Total Percent Discards
Weighting	30%	30%	20%	20%
"Excellent" Score Range	0-70%	0-70%	under 5%	under 5%
"Good" Score Range	70-80%	70-80%	5-7%	5-7%
"Warning" Score Range	80-90%	80-90%	7-10%	7-10%
"Critical" Score Range	>90%	>90%	over 10%	over 10%

For more information about the Grade of Service metric, see Table 2-2 on page 2-6. For more information about Grade of Service as it applies to each type of report, see "Analyzing TREND Reports" on page 3-1.

TREND

6 Router ReportPack

The Router ReportPack provides comprehensive data collection and reporting to support management of router networks. It includes compound reports for capacity planning, service level reporting, forecasting, and troubleshooting. It also provides an executive summary on the performance of all routers combined and detailed reports on the performance of individual routers.

The Router ReportPack answers questions including:

- Do the router networks provide an acceptable level of service to customers or end users?
- Are there overloaded routers? If so, are there under-utilized routers that can help balance the load?
- Are performance problems due to high levels of discards, errors, excessive utilization, or other causes?
- What are the top ten routers contributing to the poor health of the network?
- Which routers are likely to degrade service in the near future unless you take preventive action?

Each report consists of a combination of charts, graphs, and tables. To produce the reports, the Router ReportPack discovers all routers, collects their volume and discard data, and then processes the data automatically.

Report Descriptions

The following sections described each report in the Router ReportPack in detail.

Executive Summary Report

The Executive Summary consists of charts and graphs that show executives and high-level managers how corporate router networks are performing in terms of volume.

Note: **The TRENDweb Administrator can configure the Router ReportPack installation, so that the Executive Summary reports will only display data for devices belonging to a specific customer. In Figure 3-1, the report displays information for only the devices belonging to the Acme company. For more information about this feature, see** "Displaying Data for Specific Customer Devices in Reports" on page 1-20.

The Executive Summary report comprises the following components:

- The **Router Inventory** summary table shows all the routers used in the enterprise, the number and types of interfaces (LAN, WAN, or other), the number of routers managed, and the volume of the routers.

Router Inventory

Thu Oct 21 1999

Vendor	WAN Interfaces	LAN Interfaces	Other Interfaces	Total Interfaces	Routers Managed	Volume in Gigabytes
All Devices	5	6	0	11	4	79.21

- The **Daily Volume** bar chart displays the total router volume for all routers over the baseline period; this gives the executive an instant view of increases or decreases in router usage over time. By default, the baseline period consists of the six weeks prior to and including the report date.

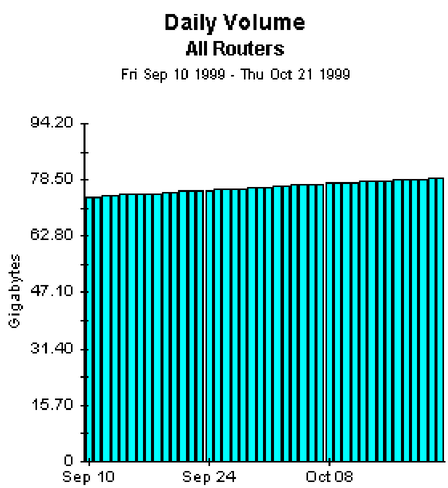

- The **Hourly Volume** bar chart displays the total router volume over the day specified and compares each hour's value to the average value for that hour over the baseline period. This chart also shows changes over time.

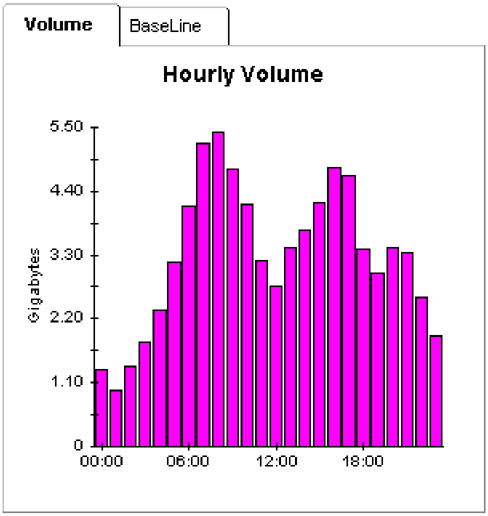

- The **Baseline Volume** line graph shows the baseline volume of the devices over the baseline period.

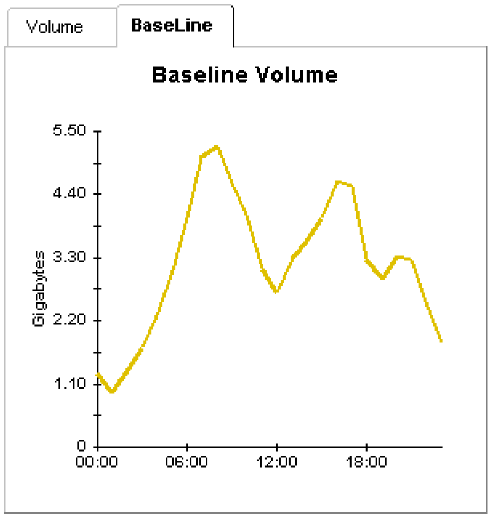

Forecast Report

The Forecast report lists all routers that are within 90 days of reaching a discard threshold. Those routers closest to the 90-day discard threshold are listed first. The calculation for the threshold is from the baseline period.

The Forecast report comprises the following components:

- The **Estimated Days to Threshold** summary table shows the number of days until the identified router exceeds the discard (5%).

Estimated Days to Threshold (DTT)
Routers Within 90 Days Of Discard Threshold

Router	Current Discards	Projected Discards 90 Days	DTT Discards	Current Pkts/sec	Projected Pkts/sec 90 Days	Current Bytes/sec	Projected Bytes/sec 90 Days
sanjose	4.85	5.01	86	3023373	4773427	133329371	191174011

Double-click on a router to update the following components with performance data for the selected router:

- The **Percent Discards** bar chart shows the percentage of discards for the selected router for each day of the baseline period.

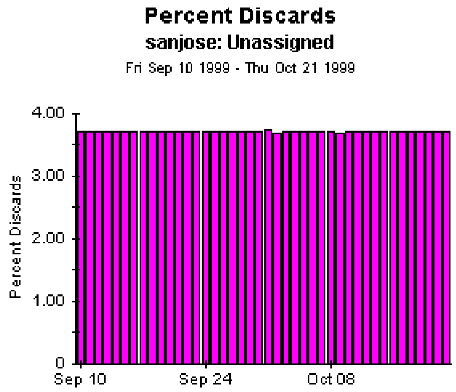

- The **Discard Rate** line graph shows the percentage of discards for the selected router for each day of the baseline period.

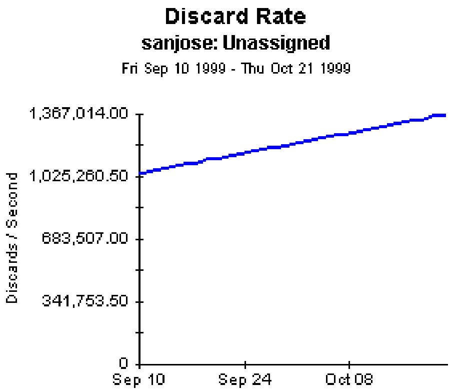

- The **Packet Rate** line graph shows the percentage of packets for the selected router for each day of the baseline period.

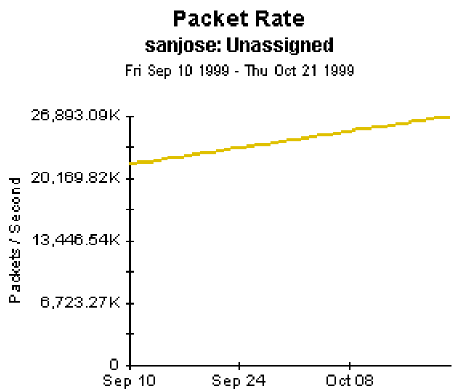

- The **Total Bytes** bar chart shows the percentage of bytes for the selected router over the baseline period.

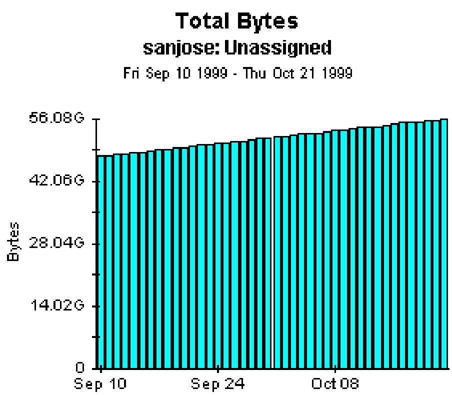

QuickView and Snapshot Reports

The QuickView report provide detailed information on key metrics for individual routers with the highest throughput. You can double-click a router in the **Router Inventory** table to display reports that show hourly volume and throughput for the selected router.

The Snapshot report provides the same information as the Quick View report, however, it allows you to select a particular router. The selected router may be different than the routers appearing in the QuickView report. For more information about selecting a router for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

The QuickView and Snapshot reports comprise the following components:

- The **Router Inventory** summary table lists the top ten routers with the highest throughput (pkts/sec) in your enterprise (QuickView) or only for routers that you select (Snapshot). For more information about selecting targets or elements for a snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

Router Inventory
Thu Oct 21 1999

Router	Vendor	WAN Interfaces	LAN Interfaces	Other Interfaces	Total Interfaces	Pkts/sec
boston	Cisco	2	1	0	3	37741747.27
tulsa	Cisco	2	1	0	3	35379888.04
lincoln	Cisco	1	N/A	0	1	33926123.71
sanjose	Bay	1	2	0	3	27538522.93
miami	Cisco	1	2	0	3	27480781.14
sanluis	Bay	1	1	0	2	27366425.47

- The **Hourly Volume** bar chart shows the volume of data flow for each hour of the report day for the selected router.

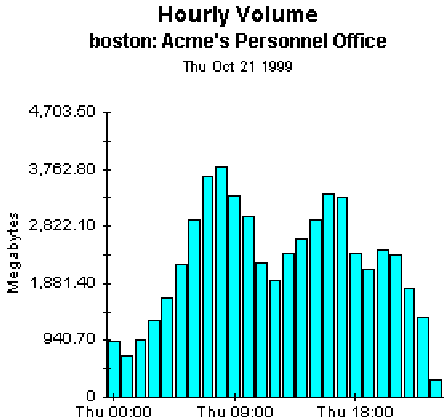

- The **Packets Per Second** line graph shows the packets/second over each hour in the report day for the selected router.

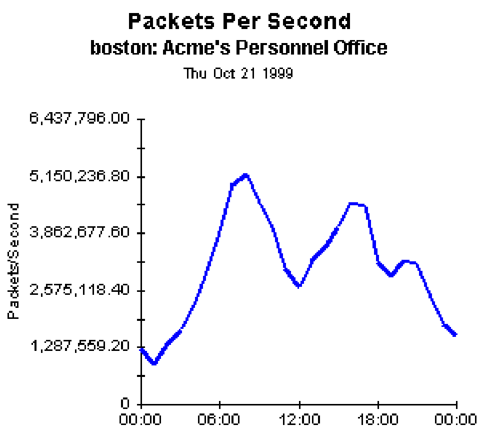

- The **Discards** bar chart shows the number of discards over each hour in the day for the selected router.

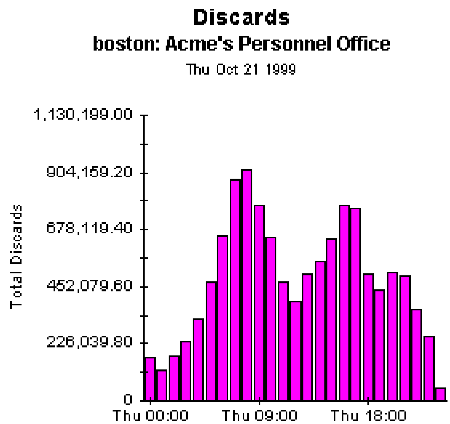

- The **Percent Discarded Packets** line graph shows the percentage of discarded packets over each hour of the report day for the selected router.

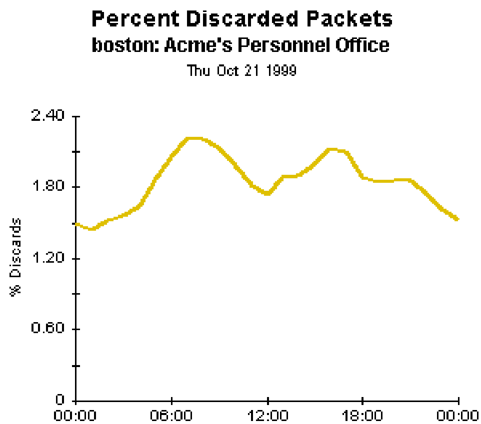

Hot Spots Report

The Hot Spots report displays the top ten routers in descending order by the total number of exceptions that exceed predetermined threshold values for detailed investigation of potential network trouble spots. Types of exceptions that contribute to the total exception count include utilization, errors, and discards exceptions.

The Hot Spots report comprises the following components:

- The **Problem Summary for the Day** table shows the routers with the highest exceptions counts ranked by total number of utilization, discard, and error exceptions.

Problem Summary for the Day
Routers With Highest Discards / Errors
Thu Oct 21 1999

Router	Vendor	Avg Percent Discards	Total Discards	Avg Percent Errors	Total Errors
miami	Cisco	3.73	32774761	3.86	33754522
sanjose	Bay	3.73	32811359	3.86	33793545
lincoln	Cisco	3.73	34531483	3.62	33417331
sanluis	Bay	3.48	28781990	3.50	27657464
houston	Cisco	2.58	7884068	1.21	3081161
cotati	Cisco	2.52	3543120	1.86	3105853
tulsa	Cisco	2.48	29244859	1.08	19032892
boston	Cisco	1.86	12021361	1.00	5195688

Double-click on a router to update the content of following charts and graphs with performance data for the selected router:

- The **Discards** bar chart shows the total number of discards for the selected router for each sample taken on the report day.

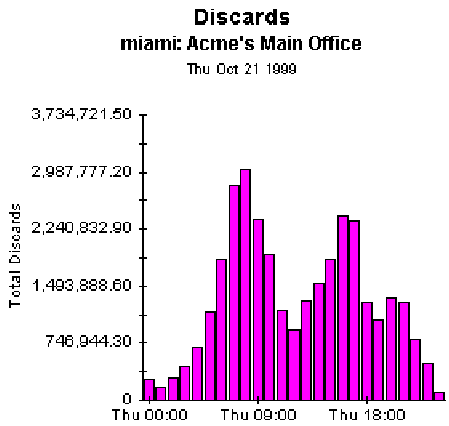

- The **Percent Discarded Packets** bar chart shows the percentage of discards for the selected router for each sample taken on the report day.

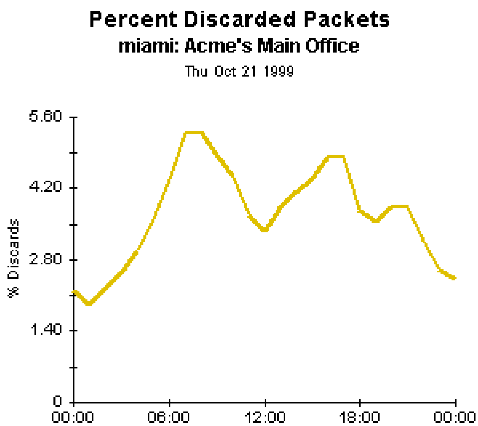

- The **Errors** bar chart shows total number of errors for the selected router for each sample taken on the preceding day.

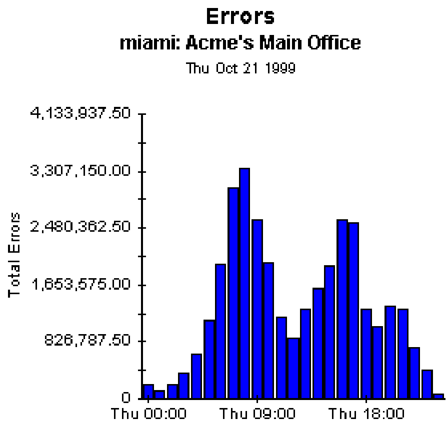

- The **Percent Errored Packets** bar chart shows the percentage of errors for the selected router for each sample taken on the report day.

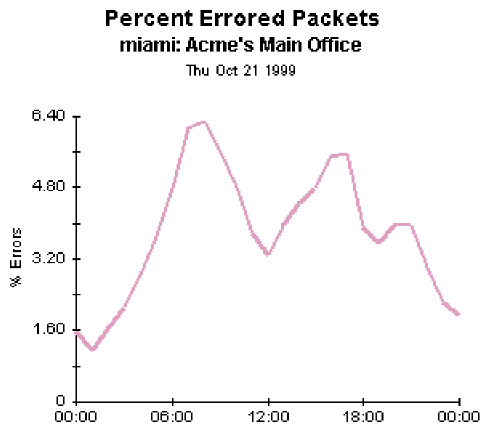

- The **Hourly Volume** bar chart shows the volume of data flow for each hour of the report day for the selected router.

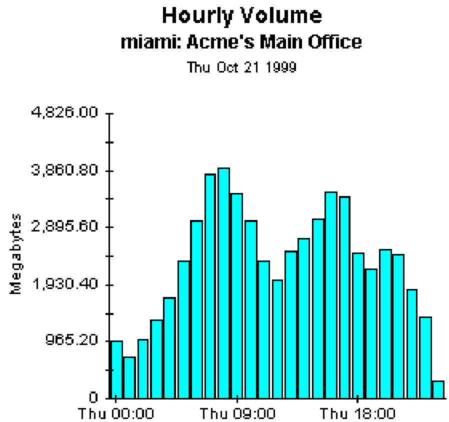

- The **Packets Per Second** line graph shows Throughput (packet/second) for the selected router for each sample period on the report day.

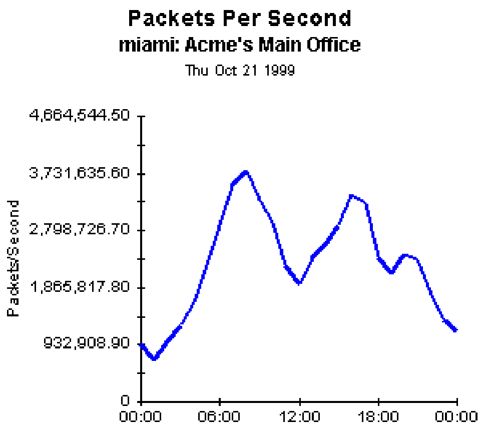

Top Ten Report

The Top Ten report lets the network manager know at-a-glance which routers have the highest throughput, highest change in throughput, highest discard rate, and highest change in discard rate. This report shows the routers with the greatest throughput that could present potential problems in the future. It ranks the top ten routers by total throughput and discard rate as well as by highest change of rank. Volume and health change charts alert users to the routers greatest changes in volume, suggesting potential capacity overloading if increasing trends continue. The volume change report also may highlight which newly installed devices and technologies are approaching a logical threshold and which may be exceeding a threshold, or which are underutilized. Drastic increases in grade of service scores suggest additional PVCs to be watched, while decreases document an effective remedy to a previously negative condition.

The Top Ten Report comprises the following components:

- The **Routers with Highest Throughput** table sorts by rank up to ten routers by the previous day's throughput in ascending order.

Routers With Highest Throughput
Thu Oct 21 1999

Router	Pkts/sec	Rank	Bytes/sec	Pct Discards
boston	37741747.27	1	1213756965.47	1.86
tulsa	35379888.04	2	947004189.12	2.48
lincoln	33926123.71	3	443368212.95	3.73
sanjose	27538522.93	4	1182195575.79	3.73
miami	27480781.14	5	1182570806.32	3.73
sanluis	27366425.47	6	853449719.02	3.48
houston	11812231.66	7	972845001.10	2.58
cotati	4015947.52	8	174540564.83	2.52

- The **Routers with Greatest Change in Throughput** table sorts by change in rank up to ten routers by previous day's throughput change in ascending order. This table may highlight which newly installed devices and technologies are approaching a logical threshold, may be exceeding a threshold, or may be impeding optimum utilization.

Routers With Greatest Change In Throughput
Thu Oct 21 1999

Router	Pkts/sec	Prev Pkts/sec	Change Rank	Rank	Prev Rank	Pct Discards
sanjose	27538522.93	27339710.09	1	4	6	3.73
tulsa	35379888.04	35477323.17	2	2	2	2.48
miami	27480781.14	27384531.32	3	5	4	3.73
boston	37741747.27	37666101.63	4	1	1	1.86
houston	11812231.66	11834926.91	5	7	7	2.58
lincoln	33926123.71	33912132.44	6	3	3	3.73
sanluis	27366425.47	27378710.53	7	6	5	3.48
cotati	4015947.52	4008552.08	8	8	8	2.52

- The **Routers with Highest Discard Rate** table sorts by rank the interfaces by previous day's discard rate in ascending order. Total throughput (packets/sec) and bytes per second are also shown.

Routers with Highest Discard Rate
Thu Oct 21 1999

Device	Pct Pkts Discarded	Rank	Prior Rank	Pkts/sec	Bytes/sec
lincoln	3.73	1	1	33926123.71	443368212.95
miami	3.73	2	2	27480781.14	1182570806.32
sanjose	3.73	3	3	27538522.93	1182195575.79
sanluis	3.48	4	4	27366425.47	853449719.02
houston	2.58	5	5	11812231.66	972845001.10
cotati	2.52	6	6	4015947.52	174540564.83
tulsa	2.48	7	7	35379888.04	947004189.12
boston	1.86	8	8	37741747.27	1213756965.47

- The **Routers with Greatest Change in Discard Rate** sorts by change in rank the interfaces by the previous day's discard rate change in descending order. Drastic increases in discard rate may suggest routers that should be watched, while decreases document an effective remedy to a previously negative condition.

Routers With Greatest Change In Discard Rate
Thu Oct 21 1999

Device	Pct Pkts Discarded	Previous Pct Pkts Discarded	Change Rank	Rank	Prior Rank	Pkts/sec	Bytes/sec
cotati	2.52	2.53	1	6	6	4015947.52	174540564.83
lincoln	3.73	3.74	2	1	1	33926123.71	443368212.95
boston	1.86	1.85	3	8	8	37741747.27	1213756965.47
sanluis	3.48	3.48	4	4	4	27366425.47	853449719.02
miami	3.73	3.73	5	2	2	27480781.14	1182570806.32
sanjose	3.73	3.73	6	3	3	27538522.93	1182195575.79
houston	2.58	2.58	7	5	5	11812231.66	972845001.10
tulsa	2.48	2.48	8	7	7	35379888.04	947004189.12

Service Level Management Report

The Service Level Management report let CIOs, CFOs, network managers, users and customers know if their routers are performing against contracted service levels. It reports the ten routers with lowest availability and the ten routers with the highest network response times. It also provides drill-down reports for daily availability and daily and hourly response times for the routers with lowest availability or highest network response time.

The Service Level Management Report comprises the following components:

◆ The **Routers with Lowest Availability** table lists the routers with the lowest availability for the report day.

Availability
Routers With Lowest Availability
Thu Oct 21 1999

Device	Availability
sanjose	100.00
sanluis	100.00
miami	100.00
tulsa	100.00
cotati	100.00
lincoln	100.00
boston	100.00
houston	100.00

Double-click a router in the table to update the following graphs with performance data for the selected router:

◆ The **Daily Availability** line graph compares the availability of the selected router with the Service Level Agreement over the baseline period. By default, the service level is set at 99.5%.

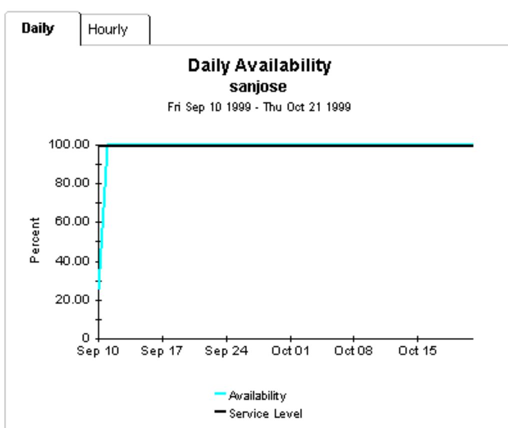

- The **Hourly Availability** line graph compares the availability of the selected router with the Service Level Agreement over the baseline period for the report day. By default, the service level is set at 99.5%.

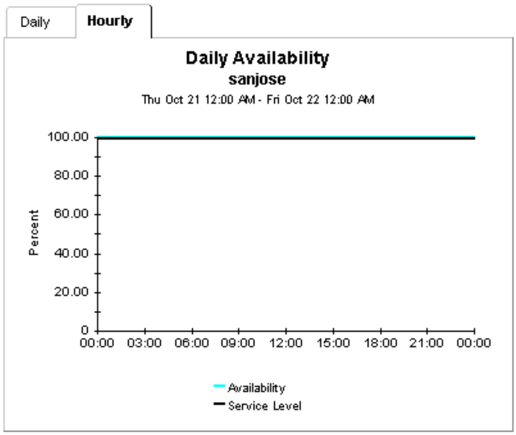

- The **Routers with Highest Response Time** table lists the routers with the highest (average) response times for the day in descending order.

Response Time

Routers With Highest Response Time

Thu Oct 21 1999

Device	Response Time
houston	558.43
miami	500.00
boston	479.52
sanluis	452.32
cotati	437.75
sanjose	386.11
tulsa	362.31
lincoln	272.95

Double-click a router in the table to update the following graphs with performance data for the selected router:

- The **Daily Response Time** line graph shows the average response time for the selected router for each day of the baseline period.

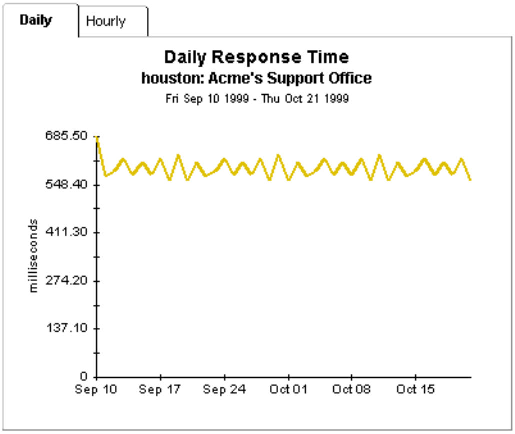

- The **Hourly Response Time** line graph shows the response time for each sample for the report day for the selected router.

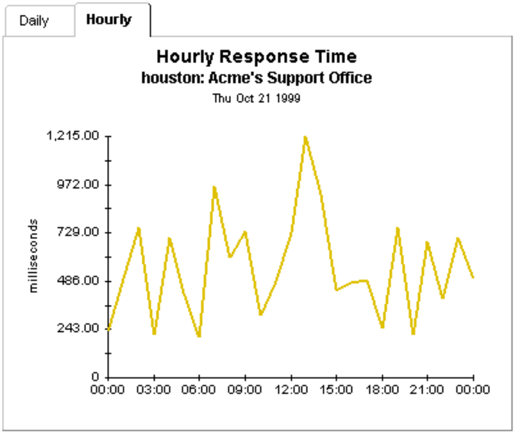

Data Source

Source data for the Router ReportPack is collected using SNMP from several MIB II tables (ifEntry, system, and ipAddrEntry).

Table 6-1 lists the variables that are polled from the MIB II ifEntry table.

Table 6-1: Variables Polled from MIB II ifEntry Table

Variable	**OID String**
ifIndex	
A unique value for each interface. Its value ranges between 1 and the value of ifNumber. The value for each interface must remain constant at least from one reinitialization of the entity's network management system to the next reinitialization.	
ifDescr	1.3.6.1.2.1.2.2.1.2
A textual string containing information about the interface. This string should include the name of the manufacturer, the product name, and the version of the hardware interface.	
ifType	1.3.6.1.2.1.2.2.1.3
The type of interface, distinguished according to the physical/link protocol(s) immediately “below” the network layer in the protocol stack.	
ifSpeed	1.3.6.1.2.1.2.2.1.5
An estimate of the interface's current bandwidth in bits per second. For interfaces which do not vary in bandwidth or for those where no accurate estimation can be made, this object should contain the nominal bandwidth.	
ifPhysAddress	1.3.6.1.2.1.2.2.1.6

(1 of 3)

Table 6-1: Variables Polled from MIB II ifEntry Table

Variable	**OID String**
The interface's address at the protocol layer immediately “below” the network layer in the protocol stack. For interfaces which do not have such an address (e.g., a serial line), this object should contain an octet string of zero length.	
ifAdminStatus	1.3.6.1.2.1.2.2.1.7
The desired state of the interface. The testing (3) state indicates that no operational packets can be passed.	
ifOperStatus	1.3.6.1.2.1.2.2.1.8
The current operational state of the interface. The testing (3) state indicates that no operational packets can be passed.	
ifLastChange	1.3.6.1.2.1.2.2.1.9
The value of sysUpTime at the time the interface entered its current operational state. If the current state was entered prior to the last reinitialization of the local network management subsystem, then this object contains a zero value.	
ifInOctets	1.3.6.1.2.1.2.2.1.10
The total number of octets received on the interface, including framing characters.	
ifInDiscards	1.3.6.1.2.1.2.2.1.13
The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space.	
ifInErrors	1.3.6.1.2.1.2.2.1.14

(2 of 3)

Table 6-1: Variables Polled from MIB II ifEntry Table

Variable	**OID String**
The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol.	
ifOutOctets	1.3.6.1.2.1.2.2.1.16
The total number of octets transmitted out of the interface, including framing characters.	
ifOutDiscards	1.3.6.1.2.1.2.2.1.19
The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space.	
ifOutErrors	1.3.6.1.2.1.2.2.1.20
The number of outbound packets that could not be transmitted because of errors.	

(3 of 3)

Table 6-2 lists the variables polled from the MIB II system table.

Table 6-2: Variables Polled from MIB II System Table

Variable	**OID String**
sysDescr	1.3.6.1.2.1.1.1
A textual description of the entity. This value should include the full name and version identification of the system’s hardware type, software operating system, and networking software. It is mandatory that this only contain printable ASCII characters.	

(1 of 3)

Table 6-2: Variables Polled from MIB II System Table

Variable	**OID String**
sysObjectId	1.3.6.1.2.1.1.2
The vendor’s authoritative identification of the network management subsystem contained in the entity. This value is allocated within the entity. This value is allocated within the SMI enterprises subtree (1.3.6.1.4.1) and provides an easy and unambiguous means for determining ’what kind of box’ is being managed. For example, if vendor ’Flintstones, Inc.’ was assigned the subtree 1.3.6.1.4.1.4242, it could assign the identifier 1.3.6.1.4.1.4242.1.1 to its ’Fred Router’.	
sysUptime	1.3.6.1.2.1.1.3
The time (in hundredths of a second) since the network management portion of the system was last reinitialized.	
sysContact	1.3.6.1.2.1.1.4
The textual identification of the contact person for this managed node, together with information on how to contact this person.	
sysName	1.3.6.1.2.1.1.5
An administratively-assigned name for this managed node. By convention, this is the node’s fully-qualified domain name.	
sysLocation	1.3.6.1.2.1.1.6
The physical location of this node (for example, ’telephone closet, 3rd floor’).	
sysServices	1.3.6.1.2.1.1.7

(2 of 3)

Table 6-2: Variables Polled from MIB II System Table

Variable	OID String
A value which indicates the set of services that this entity primarily offers. The value is a sum. This sum initially takes the value zero. Then for each layer, L, in the range 1 through 7, that this node performs transactions for, 2 raised to (L-1) is added to the sum. For example, a node which performs primarily routing functions would have a value of 4 (2^(3-1)). In contrast, a node which is a host offering application services would have a value of 72 (2^(4-1) + 2^(7-1)). Note that in the context of the Internet suite of protocols, values should be calculated accordingly: 1) layer functionality 2) physical (for example, repeaters) 3) datalink/subnetwork (for example, bridges) 4) internet (for example, IP gateways) 7) end-to-end (for example, IP hosts) 7 applications (for example, mail relays). For systems including OSI protocols, layers 5 and 6 may also be counted.	

(3 of 3)

Table 6-3 lists the variables polled from the MIB II ipAddrEntry table.

Table 6-3: Variables Polled from MIB II ipAddrEntry Table

Variable	OID String	Base Table Column Name
ipAdEntAddr	1.3.6.1.2.1.4.20.1.1	ipadentaddr009
The IP address to which this entry's addressing information pertains.		
ipAdEntIfIndex	1.3.6.1.2.1.4.20.1.2	ipadentifindex010
The index value which uniquely identifies the interface to which this entry is applicable. The interface identified by a particular value of this index is the same interface as identified by the same value of ifIndex.		
ipAdEntNetMask	1.3.6.1.2.1.4.20.1.3	ipadentnetmask011
The subnet mask associated with the IP address of this entry. The value of the mask is an IP address with all the network bits set to 1 and all the host's bits set to 0.		
ipAdEntBcastAddr	1.3.6.1.2.1.4.20.1.4	padentbcastaddr012
The value of the least-significant bit in the IP broadcast address used for sending datagrams on the (logical) interface associated with the IP address of this entry. For example, when the Internet standard all-ones broadcast address is used, the value will be 1. This value applies to both the subnet and network broadcasts addresses used by the entity on this (logical) interface.		
ipAdEntReasmMaxSize	1.3.6.1.2.1.4.20.1.5	entreasmmaxsize013
The size of the largest IP datagram which this entity can reassemble from incoming IP fragmented datagrams received on this interface.		

Basic Metric Calculations

The performance metrics used in Router reports are based on the calculations listed in Table 6-4. Note that some charts modify these calculations.

Table 6-4: Router Metrics

Metric	Computed As
Discards	discard_pct= (ifindiscards021+ifoutdiscards027)/ (ifinucastpkts019+ifinnucastpkts020+ifindiscards021+ ifinerrors022+inunknownprotos023+ifoutnucastpkts026+ ifoutucastpkts025+1)*100.0
The ratio of the sum of inbound and outbound packets containing errors and those that are discarded because they come from an unknown source to the total packet traffic (inbound and outbound) at the interface plus 1. (In this computation, a negligible amount (that is, 1 packet) is added to the divisor to prevent a potential divide-by-zero condition.) The result is multiplied by 100 to yield a percentage.	
Throughput	(TOT(ifinucastpkts+ifoutucastpkts))/delta_time
The total number of ifinucastpkts plus ifoutucastpkts (packets) for all devices divided by the total delta time for the aggregation period to yield packets per second.	
Volume	ifinoctets+ifoutoctets
inoctets and outoctets (bytes) for the selected managed elements are totaled for the specified aggregation period.	

TREND

7 Bay Router ReportPack

The Bay Router ReportPack reports statistics about CPU, memory, and buffer utilization and grade of service specifically for Bay Routers. The Router ReportPack (Chapter 6) includes Bay routers in its statistics for discards, throughput, and traffic volume.

The Bay Router ReportPack answers questions including:

- Do the Bay routers provide an acceptable level of service?
- Are performance problems due to high levels of CPU, memory, or buffer utilization or other causes?
- Which Bay routers are likely to degrade service in the near future if you do not take preventive action?

Report Descriptions

The following sections describe each report in the Bay Router ReportPack in detail.

Executive Summary Report

The Executive Summary report provides CFOs, CIOs, and other managers an overview of the performance of corporate Bay routers in the critical areas of volume; and CPU, memory, and buffer utilization. Each chart shows key metrics aggregated for all Bay routers. Key indicators of performance are shown individually and combined into a Grade of Service chart to reveal Bay router health at a glance.

The Executive Summary report comprises the following components:

- The **Bay Router Inventory** summary table shows the number of LAN, WAN, and Other interfaces; total interfaces (LAN + WAN +Other); number of Bay routers that are managed; total volume for all Bay router LAN, WAN, and Other interfaces.

Bay Router Inventory

Sat Jan 01 2000 - Sat Jan 01 2000

Vendor	WAN Interfaces	LAN Interfaces	Other Interfaces	Total Interfaces	Routers Managed	Volume in Terabytes
Bay	1	2	0	3	1	26.43

- The **Hourly Volume** bar chart shows the total volume for all Bay routers for each hour (yesterday).

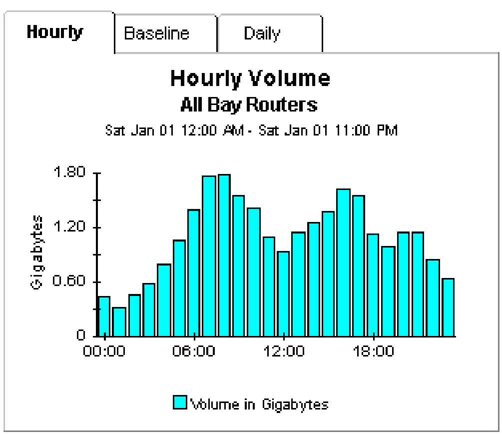

- The **Hourly Baseline** line chart shows the average volume for each hour of the day over the baseline period. You can compare total volume for each hour of the report day with the baseline for that hour for the report period.

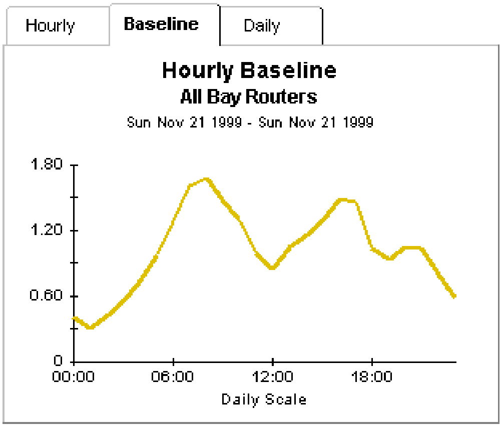

- The **Daily Volume** bar chart shows the total volume for all interfaces on all Bay routers over each day of the baseline period up to and including the report day.

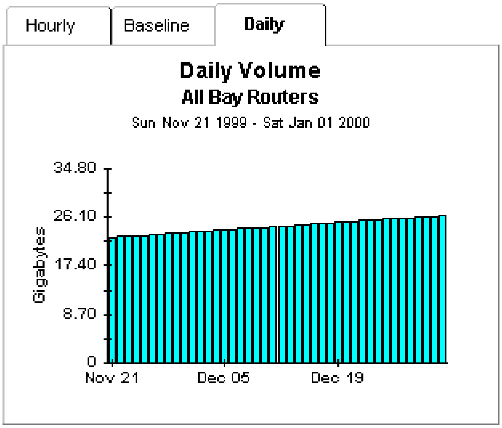

- **Hourly Grade of Service** stacked bar chart provides an overview of Bay router performance by supplying weighted values for CPU, memory, and

buffer utilization for each hour of the report day. Table 7-3 describes the GOS weighting factors used in this chart.

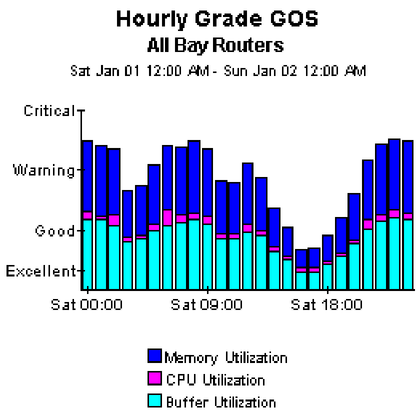

- The **Hourly CPU Utilization** line graph shows average and 95th percentile CPU utilization for each hour of the report day.

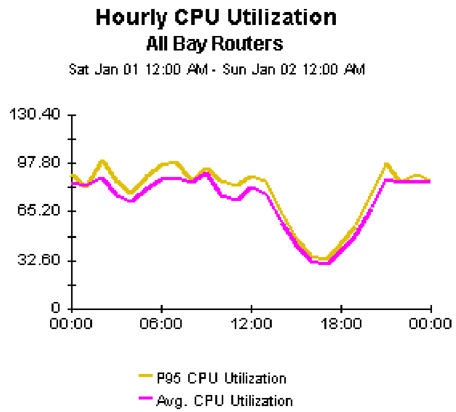

- The **Hourly Buffer Utilization** line graph shows average and 95th percentile buffer utilization for each hour of the report day.

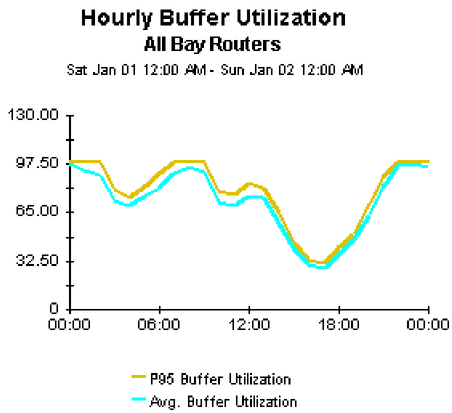

- The **Total Number of Exceptions** stacked bar graph shows the number of times a Bay router exceeded a CPU, memory, or buffer utilization threshold during each hour of the report day.

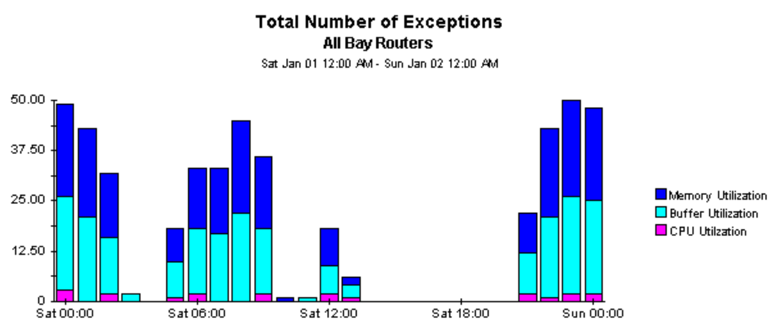

7 Bay Router Report-Pack

Forecast Report

The Forecast report lists Bay routers that are within 90 days of reaching a CPU (85%), memory (80%), or buffer utilization (80%) threshold, indicating where near-term correction can prevent service degradation. The calculation for the threshold is from the baseline period. Those Bay routers closest to the 90-day threshold are listed first. The forecast data is generated using the daily 95th percentile values for CPU, memory, and buffer utilization. Drill-down charts show CPU, memory, and buffer utilization; grade of service; discards; and total bytes over the baseline period.

Bay routers are ranked in ascending order by CPU, memory, and buffer utilization DTT.

The Forecast report comprises the following components:

- The **Estimated Days to Threshold** summary table provides the number of days until CPU, memory, and buffer utilization for the router reaches the defined thresholds and the projected 95th percentile CPU, memory, and buffer utilization for the router in 90 days.

Estimated Days To Threshold

Bay Routers Within 90 Days of Threshold

(Thresholds: CPU Utilization = 80%; Memory Util. = 75%; Buffer Util. = 75%

Router	Slot	DTT CPU Util Forecast	DTT Memory Util Forecast	DTT Buffer Util Forecast	90 Day CPU Util Forecast	90 Day Memory Util Forecast	90 Day Buffer Util Forecast
sanjose	5	N/A	-1000.00	N/A	92.23	100.00	98.75
sanluis	1	N/A	-1000.00	N/A	91.96	100.00	98.75
sanjose	3	N/A	-1000.00	N/A	91.39	100.00	98.75
sanjose	1	-1000.00	-541.00	-1000.00	94.27	98.45	95.87
sanluis	3	-1000.00	-1000.00	-1000.00	85.98	91.90	85.25

Double-click on a Bay Router in the summary table to update the content of following components with performance information for the selected router:

- The **Bay Router Daily GOS** stacked bar chart shows weighted values for CPU, memory, and buffer utilization for each day of the baseline period

for the selected router. See Table 7-3 for a description of the GOS weighting factors.

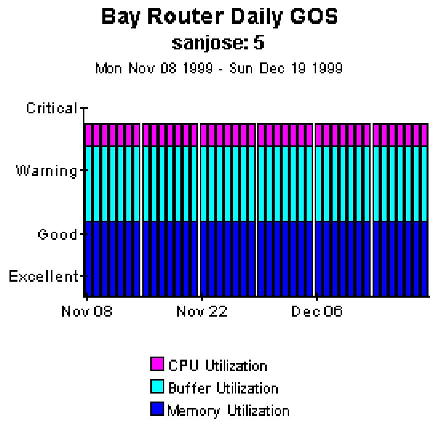

- The **Bay Router Daily CPU Utilization** line graph shows the average and 95th percentile CPU utilization for the selected router for each day of the baseline period.

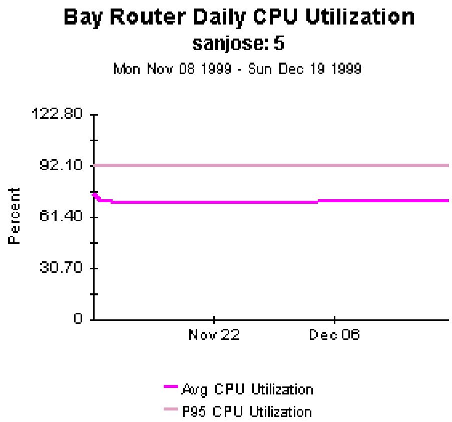

- The **Bay Router Daily Buffer Utilization** line graph shows the average and 95th percentile buffer utilization for the selected router for each day of the baseline period.

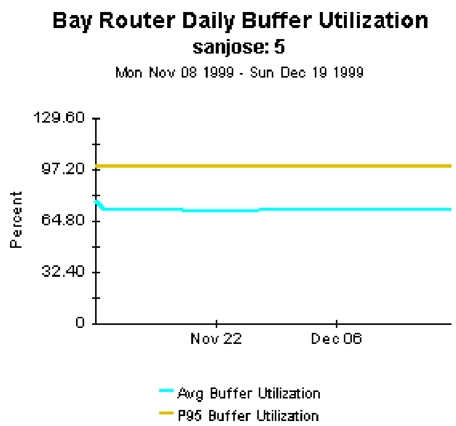

- The **Bay Router Daily Memory Utilization** line graph shows the average and 95th percentile memory utilization for the selected router for each day of the baseline period.

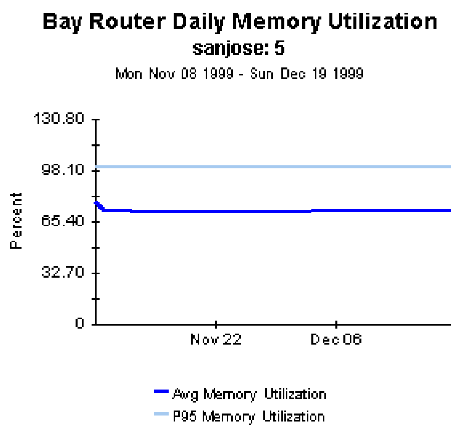

QuickView and Snapshot Reports

The QuickView report provides the network management staff with detailed information on key metrics for individual Bay routers with the highest GOS scores. The Snapshot report provides the same information as the QuickView report for individual Bay routers that you select from a pick list for inclusion in the report. For more information about selecting a Bay router for a Snapshot report, see“Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

The QuickView and Snapshot reports comprise the following components:

- The **Daily Grade of Service** stacked bar graph shows the GOS scores for the selected Bay routers. The chart shows weighted values for CPU, memory, and buffer utilization for the report day. See Table 7-3 for a description of the GOS calculation formulas.

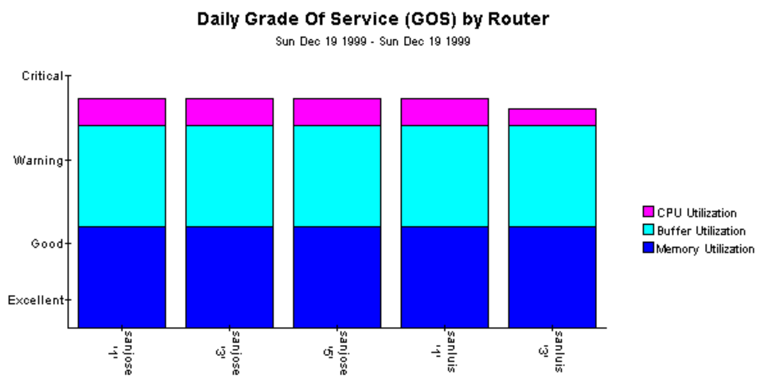

Double click on a Bay router to update the following components with performance data for the selected interface:

7 Bay Router Report-Pack

- The **Hourly Volume** bar graph shows the volume of data for each sample period of the report day for the selected Bay router.

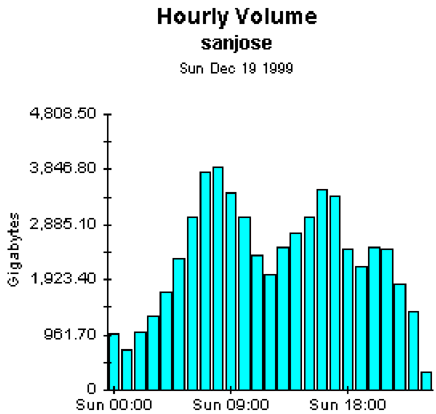

- The **CPU Utilization Frequency Distribution** pie chart shows the percentage of CPU utilization values, computed from each sample taken during the report day, that fall within the defined utilization buckets for the selected Bay router.

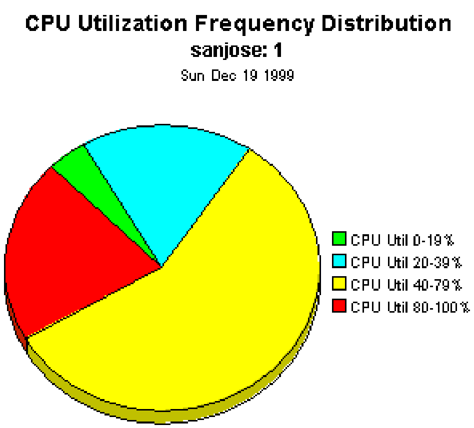

- The **Memory Utilization Frequency Distribution** pie chart shows the percentage of memory utilization values, computed from each sample taken during the report day, that fall within the defined utilization buckets for the selected Bay router.

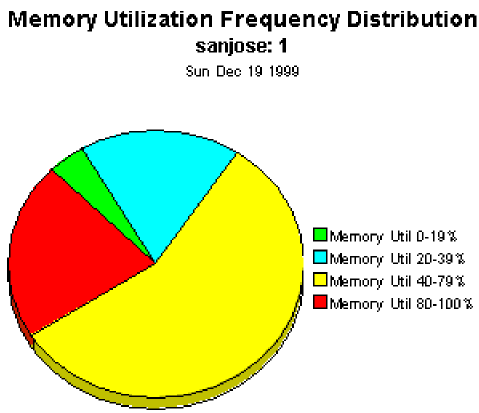

- The **Buffer Utilization Frequency Distribution** pie chart shows the percentage of buffer utilization values, computed from each sample taken during the report day, that fall within the defined utilization buckets for the selected Bay router.

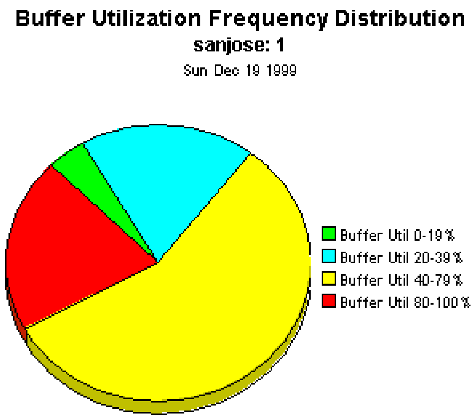

- The **Percent Errors** line graph shows the percent errors for the selected Bay router over the samples for each hour of the report day. Table 7-2 explains how percent errors is computed.

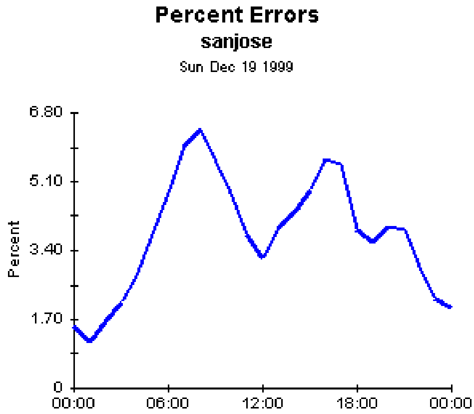

- The **Total Number of Discarded Packets** line graph shows the total number of discarded packets (ifindiscards+ifoutdiscards) for the selected Bay router over each hour of the report day.

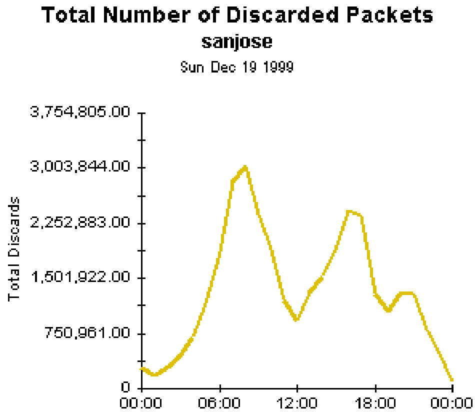

Near Real Time QuickView and Snapshot Reports

The Near Real Time--QuickView report provides router performance statistics up to the last SNMP poll. The report does not rely on nightly summaries; therefore, it provides instant reporting on collected data. Network managers and analysts can go to this report to select a router from a table, and assess its status in the following performance areas: CPU utilization, buffer utilization, and memory utilization.

Alternatively, the Snapshot report shows the same information, however, you select routers you want to view from a pick list when you invoke the report. For more information about selecting a router for a Snapshot report, see "Selecting Deferred Targets or Elements for a Snapshot Report" on page 1-27.

Note: For an in-depth analysis of how to interpret the Bay Router Near Real Time QuickView and Snapshot reports and for chart illustrations, see "Near Real Time QuickView and Snapshot Reports" on page 3-27**.**

The QuickView and Snapshot reports comprise the following charts:

- The **Router Selection List** table shows each router and its performance statistics based on the data polled for the last 6 hours. The statistic shown for each metric is the average of all the data collected for that statistic over the 6-hour period.

  Double-click on a device in the table to update the content of the following drill-down components with performance data for the device:

  - The **CPU Utilization** line graph shows the CPU utilization for each sample taken during the report period for the selected router.
  - The **Buffer Utilization** line graph shows the buffer utilization for each sample taken during the report period for the selected router.
  - The **Memory Utilization** line graph shows the memory utilization for each sample taken during the report period for the selected router.

Capacity Planning Report

The Capacity Planning report shows the most overutilized and underutilized Bay routers that are within 90 days of reaching a CPU (85%), memory (80%), or buffer utilization (80%) threshold. It allows the CIO and network manager to estimate utilization levels in 30, 60, or 90 days. You can double-click on any interface in the Overutilized Bay routers or Underutilized Bay routers summary tables to display corresponding drill-down charts that detail daily data for Grade of Service; and CPU, memory, and buffer utilization for the baseline period.

The Capacity Planning report comprises the following components:

- The **Overutilized Bay Routers** summary table shows the Bay routers that are projected to be overutilized within the next 90 days ranked in descending order by the 90-day projected CPU, memory, or buffer utilization value. All calculations are based on 95th percentile utilization values and are displayed as percentages in this table. Table 7-2 describes how CPU, memory, and buffer utilization is computed. *Overutilized* means that the router has exceeded the CPU(85%), memory (80%), or buffer (80%) utilization threshold defined.

Overutilized Bay Routers

Projected to Exceed Utilization Threshold Within 90 Days

(Thresholds: CPU Utilization = 80%; Memory Util. = 75%; Buffer Util. = 75%

Router	Slot	30 Day CPU Util Forecast	60 Day CPU Util Forecast	90 Day CPU Util Forecast	DTT CPU Util Forecast	30 Day Memory Util Forecast	60 Day Memory Util Forecast	90 Day Util F
sanjose	6	97.17	97.15	97.13	N/A	100.00	100.00	10
sanjose	2	96.58	96.54	96.50	N/A	100.00	100.00	10
sanjose	5	92.32	92.27	92.23	N/A	100.00	100.00	10
sanjose	3	91.44	91.42	91.39	N/A	100.00	100.00	10
sanjose	1	93.86	94.07	94.27	-1000.00	96.70	97.58	9
sanjose	4	85.79	86.01	86.23	-78.00	92.15	92.28	9
sanluis	3	9.69	9.69	9.69	N/A	91.51	91.71	9
sanluis	1	7.80	7.80	7.80	N/A	100.00	100.00	10

Double-click on a Bay router in this table to update the following components with performance data for the selected router:

- The **Bay Router Daily Grade of Service** stacked bar graph shows GOS scores for CPU, memory, and buffer utilization over the baseline period

for the selected router. See Table 7-3 for a description of the GOS weighting factors.

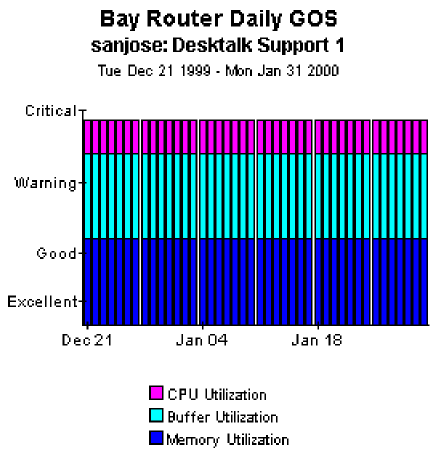

- The **Bay Router Daily CPU Utilization** line graph shows the average and 95th percentile CPU utilization over each day of the baseline period for the selected router.

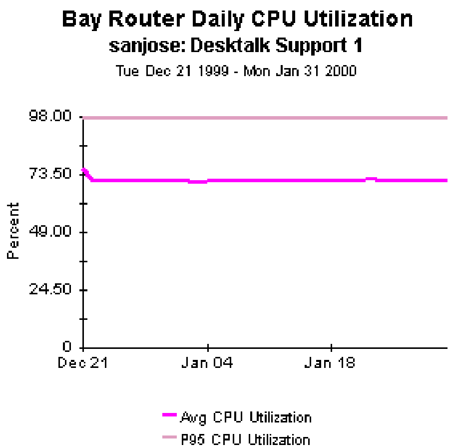

- The **Bay Router Daily Buffer Utilization** line graph shows the average and 95th percentile buffer utilization over each day of the baseline period for the selected router.

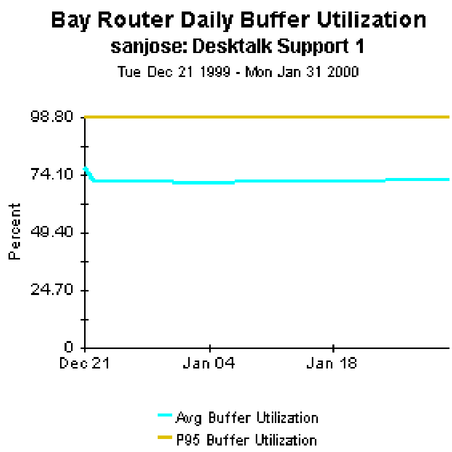

- The **Bay Router Daily Memory Utilization** line graph shows the average and 95th percentile memory utilization over each day of the baseline period for the selected router.

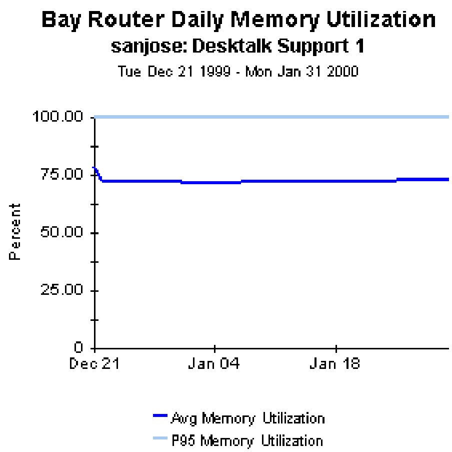

- The **Underutilized Bay Routers** summary table shows the Bay routers that projected to be the least utilized within the next 90 days ranked in ascending sequence by the 90-day projected CPU, memory, or buffer utilization value. All calculations are based on 95th percentile utilization values and are displayed as percentages in this table. *Underutilized* means that the router is furthest from reaching the utilization thresholds (10%) defined.

Underutilized Bay Routers
Projected To Be Least Utilized In 90 Days
(Thresholds: CPU Util. = 10 %; Memory Util. = 10%; Buffer Util. = 10 %

Router	Slot	30 Day CPU Util Forecast	60 Day CPU Util Forecast	90 Day CPU Util Forecast	30 Day Memory Util Forecast	60 Day Memory Util Forecast	90 Day Memory Util Forecast	30 I Util
sanluis	3	9.69	9.69	9.69	91.51	91.71	91.90	
sanluis	1	7.80	7.80	7.80	100.00	100.00	100.00	
sanluis	2	8.58	8.58	8.58	32.07	32.11	32.14	

Double-click on a Bay router in this table to update the following components with performance data for the selected router:

- The **Bay Router Daily Grade of Service** stacked bar graph shows GOS scores for CPU, memory, and buffer utilization over the baseline period for the selected router. See Table 7-3 for a description of GOS weighting factors.
- The **Bay Router Daily CPU Utilization** line graph shows the average and 95th percentile CPU utilization over each day of the baseline period for the selected router.
- The **Bay Router Daily Buffer Utilization** line graph shows the average and 95th percentile buffer utilization over each day of the baseline period for the selected router.
- The **Bay Router Daily Memory Utilization** line graph shows the average and 95th percentile memory utilization over each day of the baseline period for the selected router.

Hot Spots Report

The Hot Spots report provides the network manager and other members of the network management staff a list of Bay routers that have exceeded threshold conditions. The Hot Spots report comprises the following charts:

- The **Bay Router Problem Summary** summary table displays the top ten Bay routers in descending order by total number of exceptions (CPU + memory + buffer utilization exceptions) for the report day. An *exception* occurs when an interface exceeds its defined CPU (85%), memory (80%), or buffer (80%) utilization threshold in any sample taken during the report day.

Bay Router Problem Summary

Sun Dec 19 1999 - Sun Dec 19 1999

Router	Slot	CPU Utilization Exception	Memory Utilization Exceptions	Buffer Utilization Exceptions	Total Exceptions
sanjose	3	3.00	43.00	42.00	88.00
sanjose	5	3.00	43.00	42.00	88.00
sanluis	1	7.00	40.00	37.00	84.00
sanjose	1	3.00	21.00	20.00	44.00
sanluis	3	1.00	16.00	10.00	27.00

Double-click on any Bay router listed in the table to update the following components with performance data for the selected router:

- The **Hourly Grade of Service** stacked bar chart shows weighted values for memory, buffer, and CPU utilization for the selected Bay router over each hour on the report day. Table 7-3 gives the GOS weighting factors.

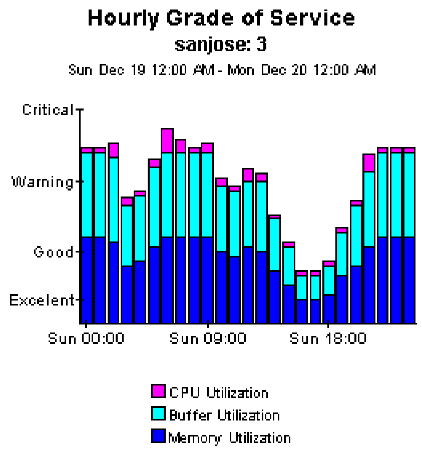

- The **Bay Router Hourly CPU Utilization** line graph shows the percentage of CPU utilization for each hour of the report day for the specified Bay router.

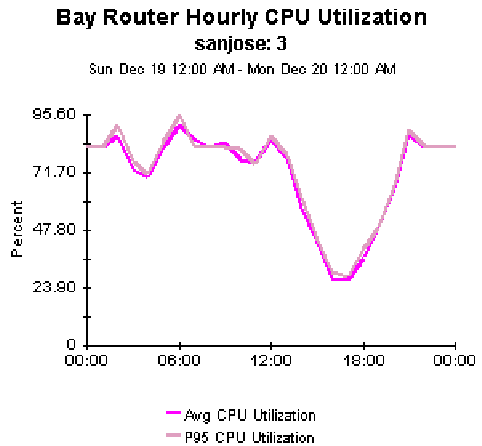

- The **Bay Router Hourly Buffer Utilization** line graph shows the average percentage of buffer utilization for each hour of the report day for the specified router.

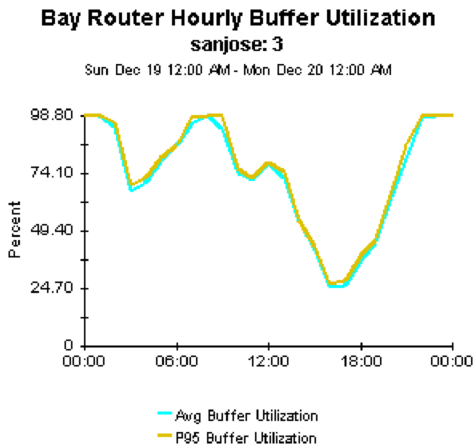

- The **Bay Router Memory Utilization** line graph shows the average percentage of memory utilization for each hour of the report day for the specified router.

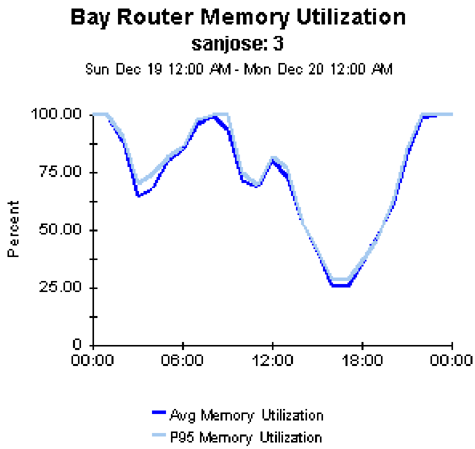

- The **Exception Detail** table shows the time the sample was taken for each CPU, buffer, and memory utilization exception and the values for the offending conditions. Table 7-2 describes how each metric is calculated.

Exception Detail Report
sanjose: 3
Sun Dec 19 12:00 AM - Mon Dec 20 12:00 AM

Time Period	Router	Slot	Description	CPU Utilization	Memory Utilization	Buffer Utilization
Sun Dec 19 12:00 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 12:15 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 12:30 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 12:45 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 01:00 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 01:15 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 01:30 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 01:45 AM	sanjose	3		82.21	100.00	98.75
Sun Dec 19 02:00 AM	sanjose	3		82.21	93.81	94.09
Sun Dec 19 02:15 AM	sanjose	3		82.21	90.17	98.95

Top Ten Report

The Top Ten report lets the network manager know at-a-glance which Bay routers produce the greatest throughput and exhibit the poorest health. The top ten Bay routers, ranked by total throughput (packets), are listed by highest rank as well as highest change of rank. Throughput and throughput-change charts alert you to the routers' greatest changes in throughput, suggesting potential capacity overloading if increasing trends continue. The throughput change report also may highlight the newly installed devices and technologies that are approaching a logical utilization threshold or are underutilized.

The Top Ten report comprises the following components:

- The **Busiest Bay Routers** summary table shows up to ten Bay routers with the highest throughput (packets) for the report day.

Busiest Bay Routers

Sun Dec 19 1999

Router	Pct Discards	Pkts/sec	Rank	Bytes/sec	Prev Rank
sanjose	3.73	27538522.93	4	1182195575.79	6
sanluis	3.48	27366425.47	6	853449719.02	5

- The **Largest Bay Router Throughput Changes** summary table ranks the selected routers in ascending sequence by change in throughput (packets). The report shows the current day's discards and traffic (in packets-per-second); previous day's traffic (in packets-per-second); and current day's rank, previous day's rank, and change in rank from the previous day. A large change in rank identifies a volatile resource, and you should investigate the reasons for the high change in throughput.

Largest Bay Routers Throughput Changes

Sun Dec 19 1999

Router	Pct Discards	Pkts/sec	Rank	Change Rank	Prev Pkts/sec	Prev Rank
sanjose	3.73	27538522.93	4	1	27339710.09	6
sanluis	3.48	27366425.47	6	7	27378710.53	5

- The **Bay Routers with the Most Percentage of Discards** summary table ranks Bay routers in descending order by the highest percentage of packet discards for the report day. Packets/second and bytes/second for the selected Bay routers are also shown.

Bay Routers with Most Percentage of Discards

Sun Dec 19 1999

Device	Pct Pkts Discarded	Pkts/sec	Rank	Bytes/sec	Prior Rank
sanjose	3.73	27538522.93	3	1182195575.79	3
sanluis	3.48	27366425.47	4	853449719.02	4

- The **Bay Routers with the Greatest Change in Percentage of Discards** summary table ranks the selected routers in ascending sequence by change in discard percentage from the previous day. A large change in rank identifies a

volatile resource, and you should investigate the reasons for the high change in discard percentage.

Bay Routers With Greatest Change in Percentage of Discards

Sun Dec 19 1999

Device	Pct Pkts Discarded	Pkts/sec	Rank	Bytes/sec	Change Rank	Previous Pct Pkts Discarded	Prior Rank
sanluis	3.48	27366425.47	4	853449719.02	4	3.48	4
sanjose	3.73	27538522.93	3	1182195575.79	6	3.73	3

Service Level Management Report

The Service Level Management report lets CIOs, CFOs, network managers, and customers know if their Bay routers are meeting contracted service levels for availability and response times by reporting Bay router availability across the enterprise over time and for individual Bay routers for the day. It reports the ten Bay routers with lowest availability for the day and the ten routers with the highest network response times. It also provides drill-down reports for daily availability over the baseline period and daily and hourly response times for the selected routers with lowest availability or highest network response time.

The Service Level Management report comprises the following components:

- The **Availability** table lists the Bay routers with the lowest availability for the report day.

Availability

Routers With Lowest Availability

Sun Dec 19 1999

Device	Availability
sanjose	100.00
sanluis	100.00

Double-click on a Bay router in this chart to update the following components with performance data for the selected router:

- The **Daily Availability** line graph shows the average availability and service level of the selected Bay router for each sample over the report period (42 days).

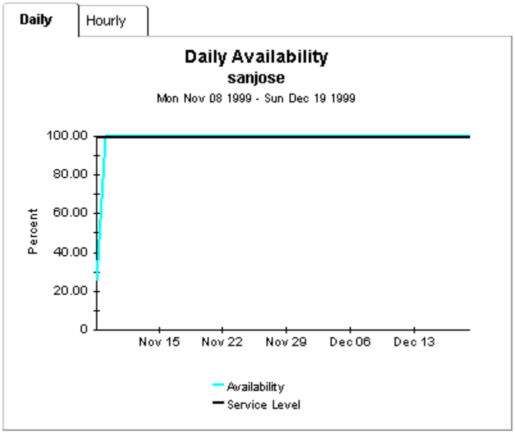

- The **Hourly Availability** line graph shows the average hourly availability and service level shows for the selected Bay router for each sample for the preceding day.

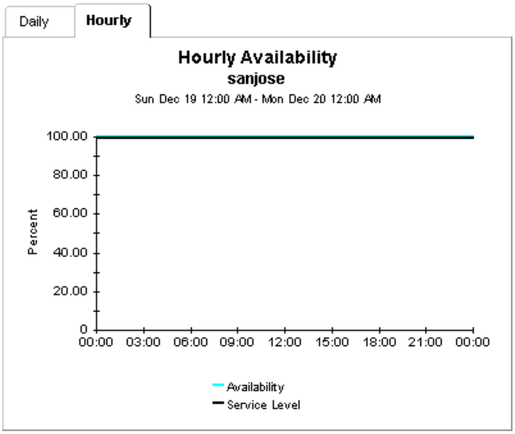

◆ The **Response Time** table lists the Bay routers with the highest (average) response times for the day.

Response Time

Routers With Highest Response Time

Sun Dec 19 1999

Device	Response Time
sanluis	452.32
sanjose	386.11

Double-click on a router in this chart to update the following components with performance data for the selected router:

◆ The **Daily Response Time** line graph shows the daily response time for the selected Bay router for each sample over the report period (42 days).

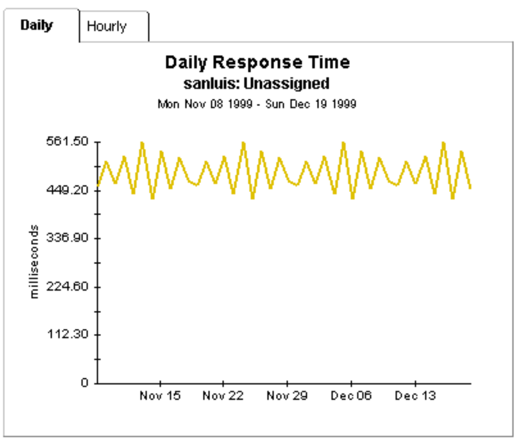

- The **Hourly Response Time** line graph shows the daily response time for the selected Bay router for each sample for the preceding day.

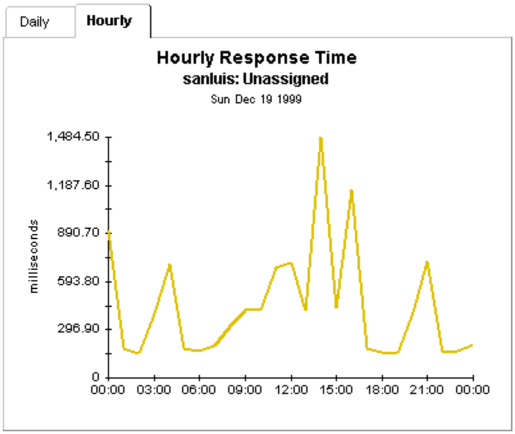

Data Source

Source data for the Bay Router ReportPack is collected from the:

- MIB II ifEntry table. These are the same as those for the Router ReportPack. See Table 7-1 for a list of variables.
- Bayrtr-resourcetotalentry table. Table 7-1 lists the variables:

Table 7-1: Bay Router Variables

Variable	OID String
wfResourceTotalSlot	1.3.6.1.4.1.18.3.3.2.5.7.1.1
The slot number is the index.	

(1 of 2)

Table 7-1: Bay Router Variables

Variable	OID String
wfResourceTotalCpuUsed	1.3.6.1.4.1.18.3.3.2.5.7.1.2
The amount of CPU used by all the entities on this slot in centiseconds since the slot started. Changed from INTEGER to Counter by DSI.	
wfResourceTotalCpuIdle	1.3.6.1.4.1.18.3.3.2.5.7.1.3
The amount of Idle CPU on this slot in centiseconds since the slot started. Changed from INTEGER to Counter by DSI.	
wfResourceTotalCpuMax	1.3.6.1.4.1.18.3.3.2.5.7.1.4
The maximum amount of CPU time possible on this slot in centiseconds since the slot started. Changed from INTEGER to Counter by DSI.	
wfResourceTotalMemoryUsed	1.3.6.1.4.1.18.3.3.2.5.7.1.5
The current number of bytes of memory used by all the entities on the slot.	
wfResourceTotalMemoryFree	1.3.6.1.4.1.18.3.3.2.5.7.1.6
The current number of bytes of free memory on the slot.	
wfResourceTotalMemoryMax	1.3.6.1.4.1.18.3.3.2.5.7.1.7
The maximum amount of Memory available in bytes on the slot.	
wfResourceTotalBuffersUsed	1.3.6.1.4.1.18.3.3.2.5.7.1.8
The current number of buffers used by all the entities on the slot.	
wfResourceTotalBuffersFree	1.3.6.1.4.1.18.3.3.2.5.7.1.9
The current number of free buffers on the slot.	
wfTotalBufersMax	1.3.6.1.4.1.18.3.3.2.5.7.1.10
The maximum number of buffers on the slot.	

(2 of 2)

Basic Metric Calculations

The performance metrics used in Bay router reports are based on the calculations listed in Table 7-2.

Table 7-2: Bay Router Metrics

Metric	Computed As ...	How Computed
CPU Utilization	((rcetotalcpuused*100)/ (urcetotalcpumax*100))* 100	The ratio of the amount of CPU time used (collected as centiseconds and multiplied by 100 to yield seconds) to the maximum amount of CPU time available (collected in centiseconds and multiplied by 100 to yield seconds) is multiplied by 100 to yield a percentage.
Memory Utilization	(totalmemoryused/ etotalmemorymax)*100	The ratio of total memory (in bytes) used to total memory available (in bytes) is multiplied by 100 to yield a percentage.
Buffer Utilization	(totalbuffersused/ totalbuffersmax)*100	The ratio of total number of buffers used to the total number of buffers available is multiplied by 100 to yield a percentage.
Volume	inoctets+outoctets	inoctets and outoctets (bytes) are totaled for the selected Bay router interface for the specified aggregation period. The source of this metric is the data collected for the LAN/WAN ReportPacks.

(1 of 2)

Table 7-2: Bay Router Metrics

Metric	Computed As ...	How Computed
Errors	(ifouterrors+ ifinerrors+ inunknownprotos)/ (ifoutucastpkts+ ifoutnucastpkts+ inunknownprotos+ ifinerrors+ ifinucastpkts+ ifinnucastpkts+ ifindiscards+1)*100.0	This metric is essentially the ratio of the sum of inbound and outbound packets containing errors and those that are discarded because they come from an unknown source to the total packet traffic (inbound and outbound) at the interface plus 1. (In this computation, a negligible amount, i.e., 1 packet, is added to the divisor to prevent a potential divide-by-zero condition.) The result is multiplied by 100 to yield a percentage. The source of this metric is the data collected for the LAN/ WAN ReportPacks.
Discards	ifindiscards+ ifoutdiscards	The number of inbound plus outbound packets that are discarded for reasons other than containing errors. The source of this metric is the data collected for the LAN/WAN ReportPacks.

(2 of 2)

Grade of Service Calculation

Table 7-3 lists how the metrics are weighted in Bay router Grade of Service (GOS) charts.

Table 7-3: Bay Router GOS Calculations

GOS Calculation	CPU Utilization	Memory Utilization	Buffer Utilization
GOS weight factor	20%	40%	40%
GOS score 1 (Excellent)	0-84%	0-20%	0-20%
GOS score 2 (Good)	85-89%	21-60%	21-60%
GOS score 4 (Warning)	90-95%	61-80%	61-80%
GOS score 5 (Critical)	95-100%	81-100%	81-100%

For more information about the Grade of Service metric, see Table 2-2 on page 2-6. For more information about Grade of Service as it applies to each type of report, see "Analyzing TREND Reports" on page 3-1.

8 Cisco Router ReportPack

The Cisco Router ReportPack provides comprehensive data collection and reporting to support management of Cisco router networks. It includes compound reports for capacity planning, service level reporting, forecasting, and troubleshooting. It also provides an executive summary on the performance of all Cisco routers combined and detailed reports on the performance of individual Cisco routers. Collectively, the reports help you identify problems on your Cisco routers, which include resource failures, not enough memory, or a router that is unable to handle the amount of traffic being generated. The reports can help you troubleshoot problems by showing what types of errors are occurring on an interface, queue drops, utilization, and protocol distribution.

The Cisco Router ReportPack answers questions including:

- Do the Cisco router networks provide an acceptable level of service to customers or end users?
- Are there overloaded Cisco routers? If so, are there underutilized Cisco routers that can help balance the load?
- Are performance problems due to high levels of discards, errors, excessive utilization, or other causes?
- What are the top ten Cisco routers contributing to the poor health of the network?

- Which Cisco routers are likely to degrade service unless you take preventive action?

Each report consists of a combination of charts, graphs, and tables. To produce the reports, the Cisco Router ReportPack discovers all Cisco routers, collects their volume and discard data, and then processes the data automatically.

Report Descriptions

The following sections describe each report in the Cisco Router ReportPack in detail.

Executive Summary Report

The Executive Summary report allows executives to review how Cisco Router networks as a whole are performing in the critical areas of volume, CPU utilization, memory utilization, and buffer utilization over time.

Note: The TRENDweb Administrator can configure the Cisco Router ReportPack installation, so that the Executive Summary reports will only display data for devices belonging to a specific customer. In Figure 3-1, the **report displays information for only the devices belonging to the Acme company. For more information about this feature, see** “Displaying Data for Specific Customer Devices in Reports” on page 1-20**.**

The Executive Summary report comprises the following components:

- The **Cisco Router Inventory** table shows all the Cisco routers used in the enterprise, the number and types of interfaces (LAN, WAN, or other), and the number of routers managed.

Cisco Router Inventory
Sun Dec 19 1999

Vendor	WAN Interfaces	LAN Interfaces	Other Interfaces	Total Interfaces	Routers Managed
Cisco	3	1	0	4	4

- The **Daily Volume** bar chart shows the total volume for all Cisco routers for each day of the baseline period giving the executive an instant view of increases and decreases in Cisco router usage over time. The default baseline period is the six weeks prior to and including the report date.

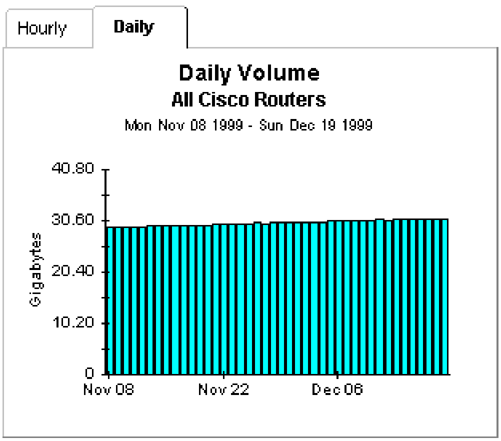

- The **Hourly Volume** bar chart shows the volume of all Cisco routers for each hour of the report day. The baseline computed for each hour over the baseline period appears as a line over the bars on the graph to compare each hour's value for that hour on that day of the week. This report reveals changes over time.

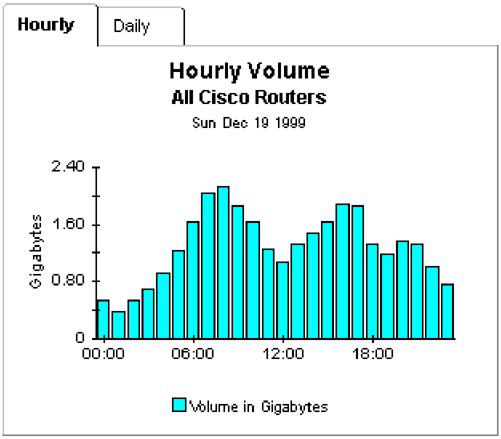

- The **Hourly Grade of Service** stacked bar chart shows grade of service scores for CPU Utilization, Physical Memory Utilization, Buffer Utilization, and Buffer Errors for each sample collected on the report day.

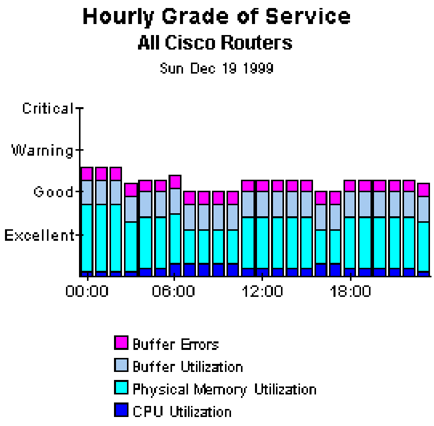

- The **Hourly CPU Utilization** line graph shows the average and 95th percentile CPU Utilization of all samples collected for each hour of the report day.

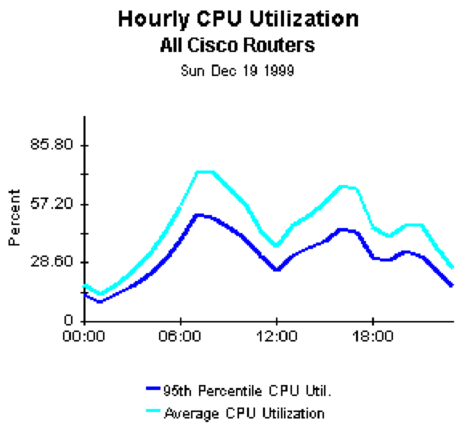

- The **Hourly Buffer Utilization** line graph shows the 95th percentile of the 95th percentile of all Buffer Utilization samples collected for each hour of the report day for each buffer size. There are five buffer sizes: huge, large, big, medium, and small.

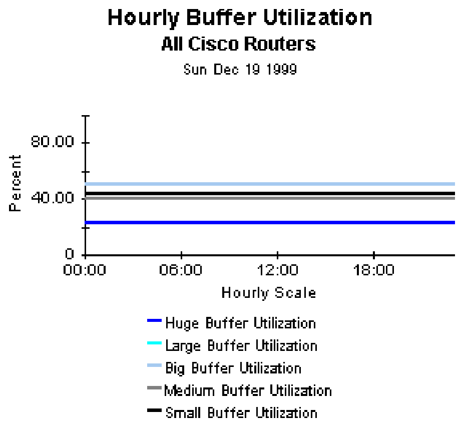

- The **Hourly Exception Counts** bar chart shows the number of exceptions for CPU Utilization, Buffer Utilization, Physical Memory Utilization, and Buffer Errors for all Cisco routers for each hour of the report day.

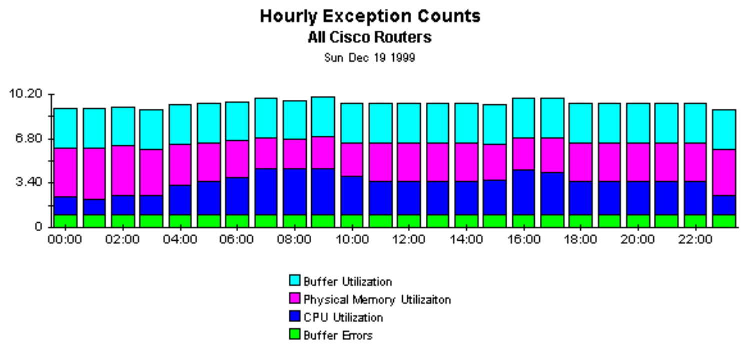

Forecast Report

The Forecast report lists all Cisco routers that are within 90 days of reaching a utilization threshold. Those Cisco routers closest to the 90-day threshold are listed first. Each managed element is hot-linked to multiple graphs for analysis over the baseline period. The default baseline period is the six weeks prior to and including the report date.

The Forecast report comprises the following components:

- The **Estimated Days to Threshold** table ranks Cisco routers that are within 90 days of reaching a 95% CPU Utilization, 95% Memory Utilization, or 75% Buffer Utilization thresholds in ascending order. For more information about the 95th percentile, days to threshold (DTT), and 90-day forecast methods of calculation, see the TRENDsum utility section in Appendix A of the *TREND User's Guide*.

Estimated Days To Threshold (DTT)
Cisco Routers Within 90 Days Of Threshold
(Thresholds: CPU Utilization=95%; Memory Util.=75%; Buffer Util.=75%)

Router	DTT CPU Utilization	90 Day CPU Util Forecast	DTT Physical Memory Util	90 Day Mem Util Forecast	DTT Buffer Util	90 Day Buffer Util Forecast
tulsa	17.00	111.49	-128.00	84.58	1000.00	44.87
boston	12.00	112.14	-128.00	84.12	1000.00	44.87
lincoln	1000.00	48.20	-545.00	90.41	N/A	59.02
houston	0.00	98.00	-1000.00	98.36	1000.00	22.90
cotati	N/A	6.99	-1000.00	98.27	-1000.00	93.78

Double-click a router in the **Estimated Days to Threshold** table to update the content of following components with performance data for the selected router:

- The **Daily Grade of Service** stacked bar chart shows grade of service scores for CPU Utilization, Memory Utilization, Buffer Utilization, and Buffer Faults for the selected Cisco router for each day of the baseline period.

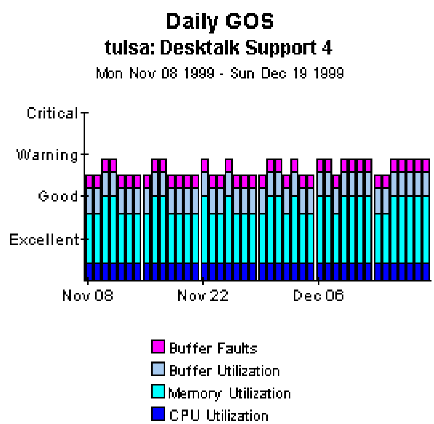

Each metric shows the 95th percentile of the value with the Grade of Service (GOS) weight applied for the selected Cisco router for each day of the baseline period. Table 8-3 describes how Grade of Service is calculated.

- The **Daily CPU Utilization** line graph shows the weighted average and 95th percentile CPU utilization for the selected Cisco router for each day of the baseline period.

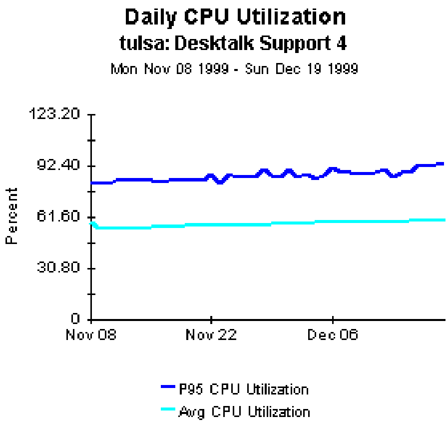

- The **Daily Memory Utilization** line graph shows the weighted average and 95th percentile memory utilization for the selected Cisco router for each day of the baseline period.

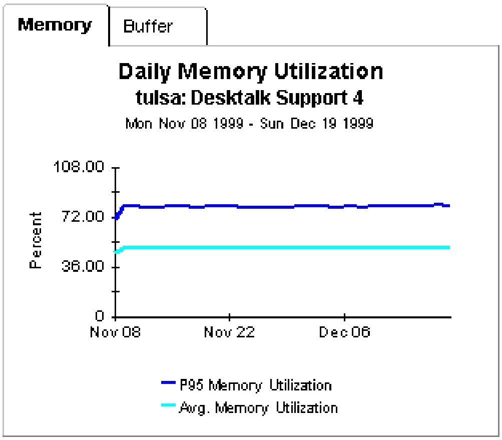

- The **Daily Buffer Utilization** line graph shows the weighted average of the buffer utilization samples collected for the selected Cisco router for each day of the baseline period for each buffer size. There are four buffer sizes: huge, large, medium, and small.

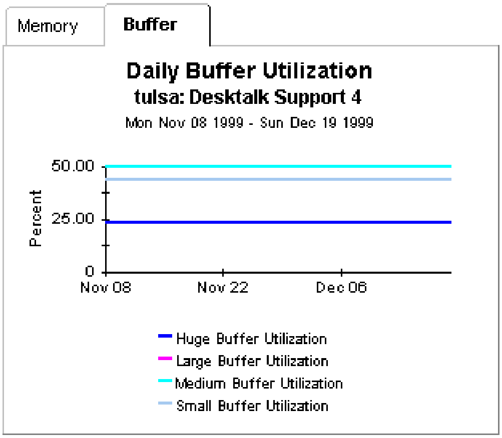

- The **Daily Total Volume** bar chart shows the total number of bytes (volume) for the selected Cisco router for each day of the baseline period. For more information about how Volume is computed, see Table 7-2.

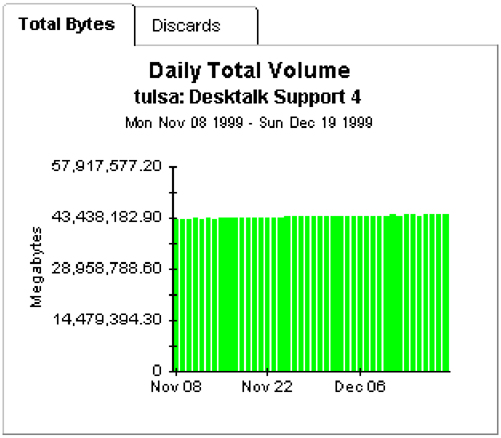

- The **Daily Percent Discards** bar chart shows the percentage of discards for the selected Cisco router for each day of the baseline period. For more information about how % Discards is computed, see Table 7-2.

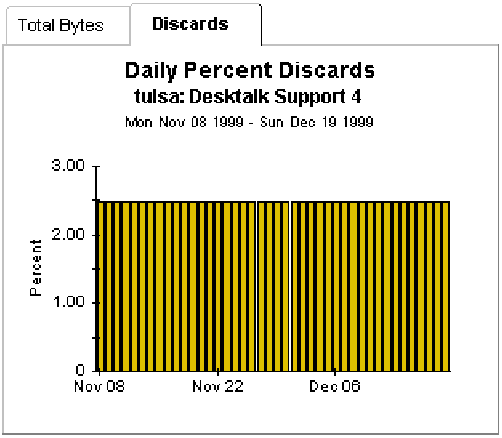

QuickView and Snapshot Reports

The QuickView report provides detailed information on individual Cisco Routers in a single report. The Snapshot report provides the same information as the Quick View report, however, it allows you to show detailed information about a particular Cisco router. The selected Cisco router may be different than the Cisco routers appearing in the QuickView report.

QuickView Report

The QuickView report comprises the following components:

- The **Grade of Service** stacked bar chart shows the Cisco routers with the highest daily GOS scores and 95th percentile of the GOS score for Memory Utilization, CPU Utilization, Buffer Utilization, and Buffer Faults. Table 8-3 describes how Grade of Service is calculated.

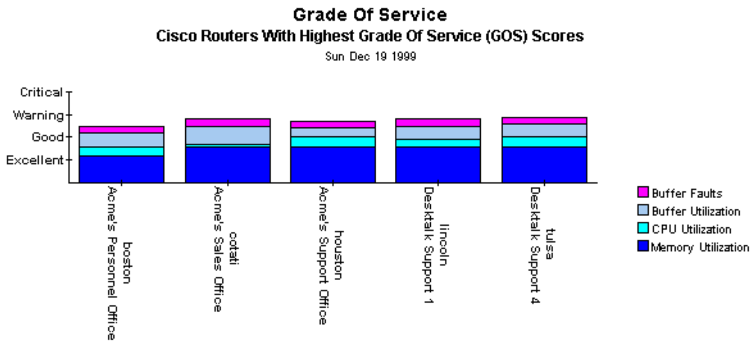

Double-click on a router in the chart to update the content of the following components with performance data for the selected router:

- The **Hourly Volume** bar chart shows the volume of the selected Cisco router for each hour of the report day. For more information about how Volume is computed, see Table 7-2.

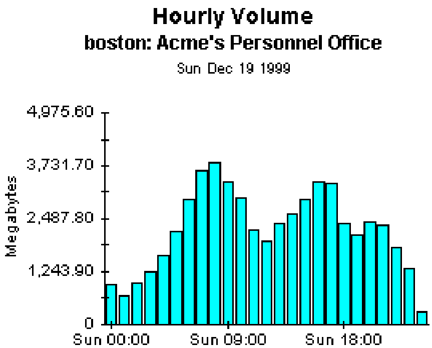

- The **CPU Utilization** pie chart shows the percentage of samples taken during the day in which the CPU Utilization metrics calculated for each sample fall in the defined range.

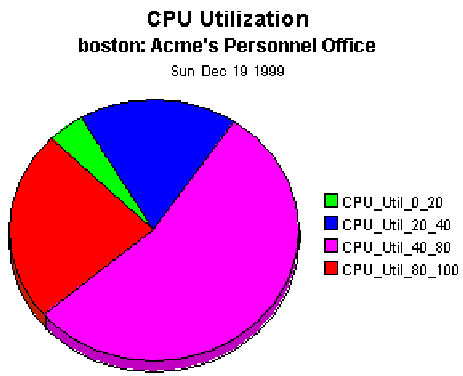

- The **Free Memory** pie chart shows the percentage of samples taken during the day in which the Free Memory Utilization metrics calculated for each sample fall in the defined range.

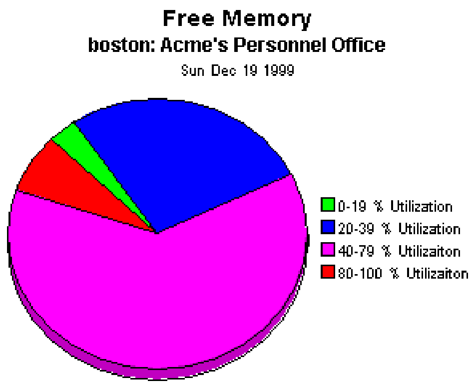

- The **Total Buffers** pie chart shows the percentage of samples taken during the day in which the Total Buffers metrics calculated for each sample fall in the defined range.

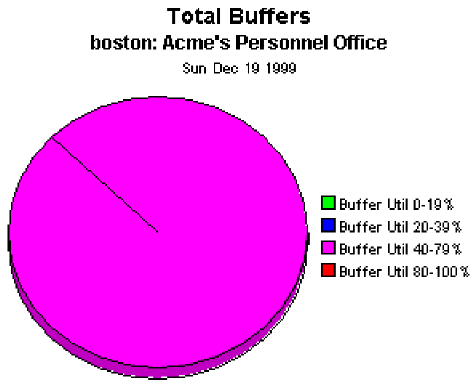

- The **Percent Errors** line graph shows the average percentage of errors for each hour during the report day. For more information about how % Errors is computed, see Table 7-2.

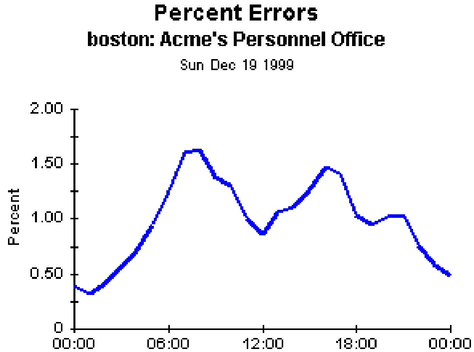

- The **Percent Discards** line graph shows the average percentage of discards for each hour during the report day. For more information about how % Discards is computed, see Table 7-2.

8 Cisco Router ReportPack

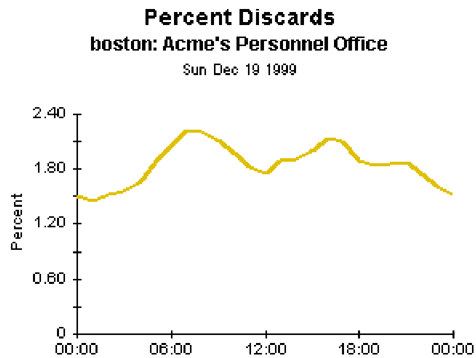

Snapshot Report

For more information about selecting a router for a Snapshot report, see "Selecting Deferred Targets or Elements for a Snapshot Report" on page 1-27.

The Snapshot report comprises the following components:

- The **Snapshot Selection** table lists the time period for which the router information was collected, the router name, any associated interfaces, and the router's CPU Utilization, Memory Utilization, and Buffer Utilization. Double-click on a router in the table to update the content of the following components with performance data for the selected router.

Snapshot Selection

Sun Dec 19 1999

Router	Interface	Description	CPU Util (95th percentile)	Memory Util (95th percentile)	Buffer Util (95th percentile)
boston	0	Acme's Personnel Office	92.00	79.96	44.87
cotati	0	Acme's Sales Office	7.00	98.07	93.78
houston	0	Acme's Support Office	98.00	97.98	22.90
lincoln	0	Desktalk Support 1	46.00	87.56	59.02
tulsa	0	Desktalk Support 4	94.00	80.62	44.87

- The **Hourly Grade of Service** bar chart shows the Grade of Service scores for Physical Memory Utilization, CPU Utilization, Buffer Utilization, and Buffer Errors for each hour of the report day. Table 8-3 describes how Grade of Service is calculated.
- The **Hourly Volume** bar chart shows the volume of the selected Cisco router for each hour of the report day. For more information about how Volume is computed, see Table 7-2.
- The **CPU Utilization**, **Free Memory**, and **Total Buffers** pie charts show the percentage of samples taken during the day in which the CPU Utilization, Free Memory Utilization, and Total Buffers metrics calculated for each sample fall in the defined range.
- The **Percent Errors** line graph shows the total percentage of errors for each hour during the day. For more information about how percent errors is computed, see Table 7-2.
- The **Percent Discards** line graph shows the percentage of discards for each hour during the day. For more information about how % Discards is computed, see Table 7-2.

Near Real Time QuickView and Snapshot Reports

The Near Real Time--QuickView report provides router performance statistics up to the last SNMP poll. The report does not rely on nightly summaries; therefore, it provides instant reporting on collected data. Network managers and analysts can go to this report to select a router from a table, and assess its status in the following performance areas: CPU utilization, buffer utilization, and memory utilization.

Alternatively, the Snapshot report shows the same information, however, you select routers you want to view from a pick list when you invoke the report. For more information about selecting a router for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

The QuickView and Snapshot reports comprise the following charts:

- The **Router Selection List** table shows each router and its performance statistics based on the data polled for the last 6 hours. Depending on the type of statistic, the value shown is either the average or maximum value of all the data collected over the 6-hour period.

Router Selection List

Select interface to see near real time reports

Router	Avg CPU Utilization	Max CPU Utilization	Small Buffer Utilization	Medium Buffer Utilization	Big Buffer Utilization	Large Buffer Utilization	Huge Buffer Utilization
houston	60.00	99.00	57.62	80.00	80.00	95.83	58.33
tulsa	39.70	68.00	56.25	50.00	50.00	50.00	76.75
boston	38.25	67.00	56.25	50.00	50.00	50.00	76.75
atlanta	34.30	58.00	75.00	75.00	73.33	68.75	86.67
lincoln	19.70	35.00	12.50	60.00	50.00	50.00	30.00
cotati	2.80	5.00	6.25	11.00	5.00	4.17	8.33

Double-click on a device in the table to update the content of the following drill-down components with performance data for the device:

- The **Average CPU Utilization** line graph shows the average CPU utilization for each sample taken during the report period for the selected router.

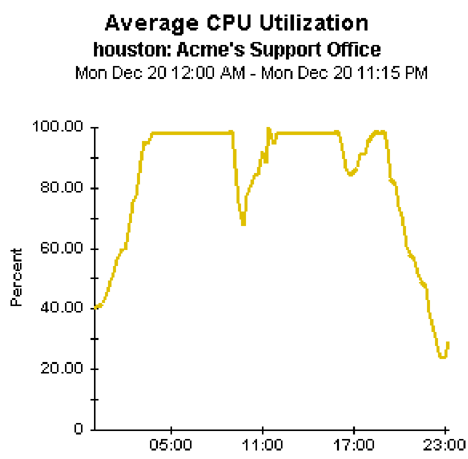

- The **Buffer Utilization** line graph shows the buffer utilization for each sample taken during the report period for the selected router for each

buffer size. There are five buffer sizes: huge, large, big, medium, and small.

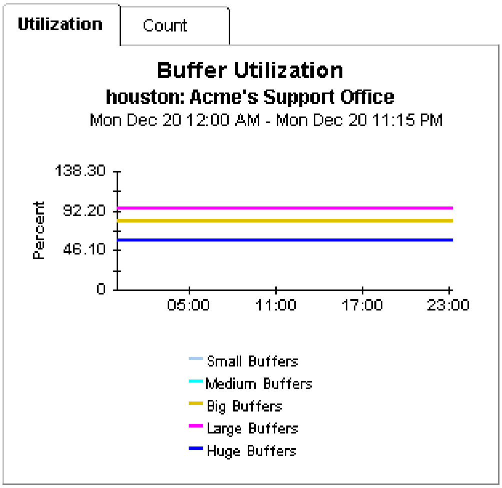

- The **Buffer Counts** line graph shows the number of buffers for each sample taken during the report period for the selected router for each buffer size. There are five buffer sizes: huge, large, big, medium, and small.

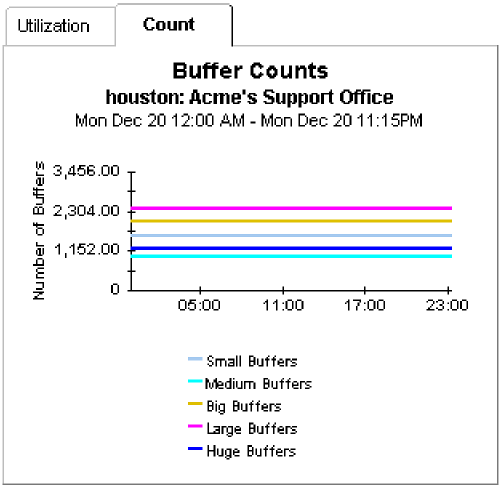

- The **Buffer Faults** line graph shows the number of buffer faults for each sample taken during the report period for the selected router.

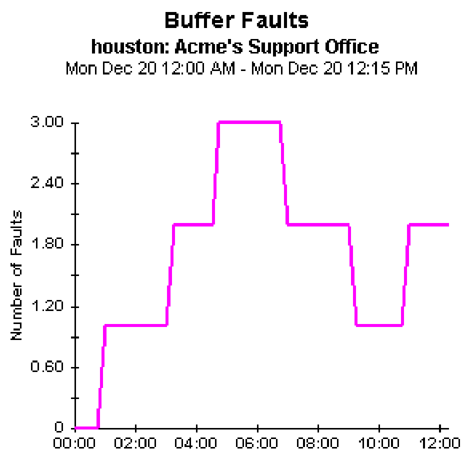

- The **Physical Memory Utilization** line graph shows the percentage of memory utilization for each sample taken during the report period for the selected router.

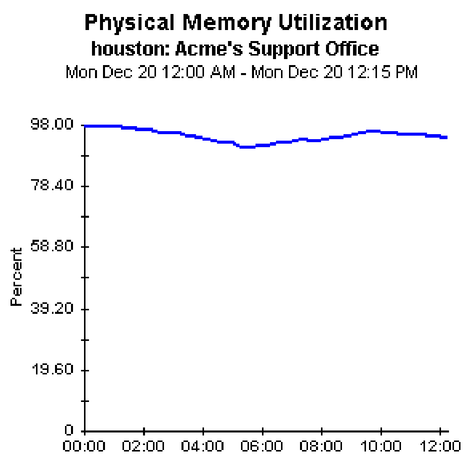

Capacity Planning Report

The Cisco Router Capacity Planning report shows the most overutilized and underutilized Cisco routers. It enables you to allocate the traffic load to achieve maximum Cisco router utilization potential.

The Capacity Planning report comprises the following components:

- The **Overutilized Cisco Routers** table ranks the Cisco routers forecasted to reach 95% CPU Utilization, 95% Memory Utilization, or 80% Buffer Utilization within the next 30, 60, or 90 days in descending order.

Overutilized Cisco Routers
Projected To Exceed Utilization Thresholds Within 90 Days
(Thresholds: CPU Utilization=95%; Memory Util.=75%; Buffer Util.=75%)

Router	DTT CPU Utilization	30 Day CPU Util Forecast	60 Day CPU Util Forecast	90 Day CPU Util Forecast	DTT Memory Util	30 Day Mem Util Forecast	60 Day Mem Util Forecast	9 U
tulsa	17.00	98.09	104.79	111.49	-128.00	81.95	83.27	
boston	12.00	99.02	105.58	112.14	-128.00	81.62	82.87	
lincoln	1000.00	47.01	47.60	48.20	-545.00	88.96	89.69	
houston	0.00	98.00	98.00	98.00	-1000.00	98.15	98.26	
cotati	N/A	7.02	7.01	6.99	-1000.00	98.10	98.18	

- The **Underutilized Cisco Routers** table ranks the Cisco routers projected to be the least utilized in terms of CPU Utilization (10%), Memory Utilization (10%), or Buffer Utilization (10%) within the next 90 days in ascending order.

Underutilized Cisco Routers
Projected To Be Least Utilized For 90 Days
(Thresholds: CPU=10%; Phys Mem=10%; Buffer=10%)

Router	DTT CPU Utilization	30 Day CPU Util Forecast	60 Day CPU Util Forecast	90 Day CPU Util Forecast	DTT Memory Util	30 Day Mem Util Forecast	60 Day Mem Util Forecast	9 U
cotati	N/A	7.02	7.01	6.99	-1000.00	98.10	98.18	

Double-click on a router in either table to update the content of the following components with performance data for the selected router:

- The **Daily GOS** chart shows grade of service scores for CPU Utilization, Memory Utilization, Buffer Utilization, and Buffer Errors for the selected Cisco router for each day of the baseline period.

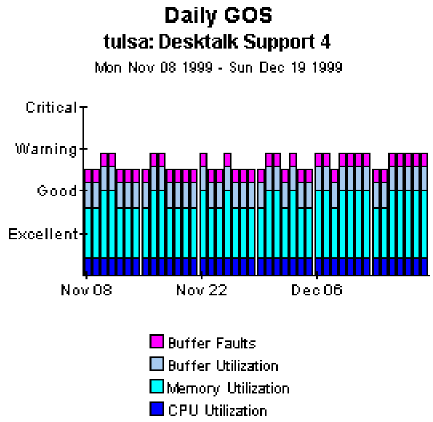

Each metric shows the 95th percentile of the value with the Grade of Service weight applied for the selected Cisco router for each day of the baseline period.

- The **Daily CPU Utilization** line graph shows the weighted average and 95th percentile CPU utilization for the selected Cisco router for each day of the baseline period.

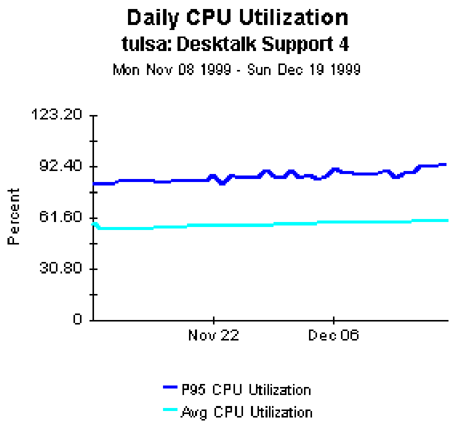

- The **Daily Memory Utilization** line graph shows the weighted average and 95th percentile memory utilization for the selected Cisco router for each day of the baseline period.

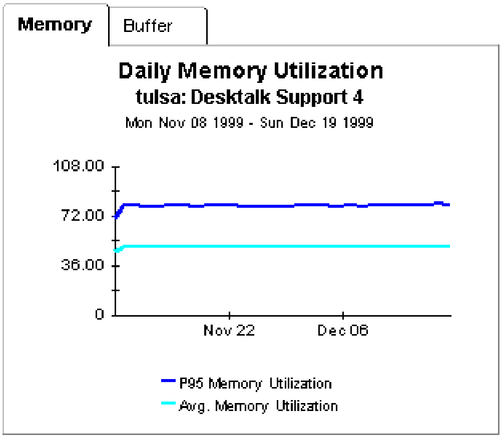

- The **Daily Buffer Utilization** graph shows the weighted average of the weighted average of the buffer utilization samples collected for the selected Cisco router for each day of the baseline period for each buffer size. There are four buffer sizes: huge, large, medium, and small.

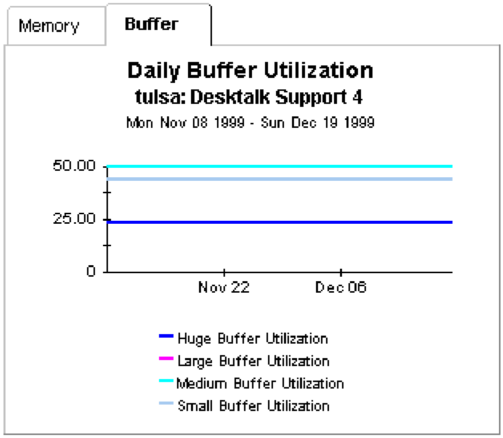

- The **Daily Total Volume** chart shows the total number of bytes (volume) for the selected Cisco router for each day of the baseline period. For more information about how Volume is computed, see Table 7-2.

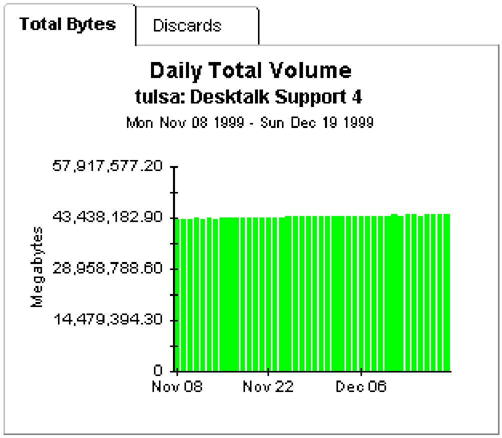

- The **Daily Percent Discards** chart shows the percentage of discards for the selected Cisco router for each day of the baseline period. For more information about how % Discards is computed, see Table 7-2.

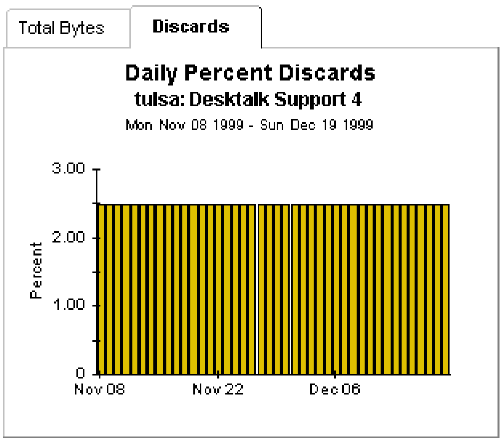

Hot Spots Report

This report lists Cisco routers that have exceeded CPU Utilization, Memory Utilization, Buffer Utilization, or Buffer Faults threshold conditions during the report day. Offending Cisco routers are ranked by total number of exceptions in the Problem Summary table. Each exception is linked to drill-down charts that display information for the following: hourly Grade of Service, CPU utilization, buffer utilization, memory utilization, percentage of discards, total buffer faults, and exception detail.

Note: For an in-depth analysis of how to interpret the Cisco Router Hot Spots report and for chart illustrations, see “Hot Spots Report” on page 3-32.

The Hot Spots report comprises the following charts:

- The **Problem Summary for the Day** table ranks the Cisco routers with the highest total CPU utilization, buffer utilization, memory utilization, and buffer fault exceptions counts by total number of exceptions (cpuutil+memutil+bufutil+bufferfaults) in descending order.

  Double-click on a router to update the content of the following components with performance data for the selected router:

  - The **Hourly Grade of Service** stacked bar chart shows the Grade of Service (GOS) ratings for CPU Utilization, Physical Memory Utilization, Buffer Utilization, and Buffer Errors for the selected Cisco router for each sample taken on the report day for the selected Cisco router. Table 8-3 defines the GOS weighting for each of these metrics.
  - The **CPU Utilization** area graph shows CPU Utilization for the selected Cisco router for each sample period on the report day.
  - The **Buffer Utilization** line graph shows Buffer Utilization for the selected Cisco router for each sample taken on the report day for each buffer size. There are five buffer sizes: huge, big, large, medium, and small.
  - The **Memory Utilization** area graph shows Memory Utilization for the selected Cisco router for each sample taken on the report day.
  - The **Percent Discards** bar chart shows the percentage of discards for the selected Cisco router for each sample taken on the report day. For more information about how the percentage of discards is calculated, see Table 7-2.
  - The **Buffer Faults** bar chart shows the total number of Buffer Faults for the selected Cisco router for each hour of the report day.
  - The **Exception Detail** table shows the number of CPU Utilization, Memory Utilization, Buffer Utilization, and Buffer Faults exception conditions occurring for each sample period on the report day for the selected Cisco router. A CPU Utilization (>95%), Memory Utilization (>75%), or Buffer Utilization (>80%) metric greater than 80% is treated as an exception and is included in the chart. For more information about how utilization and buffer faults are calculated, see Table 8-2.

Top Ten Report

The Top Ten report lets the network manager know at-a-glance which Cisco routers have the highest throughput and discard percentage. The top ten routers, ranked by throughput and discard percentage, are listed by highest rank as well as highest change of rank.

- The **Cisco Routers with Highest Throughput** table shows the ten Cisco routers with the highest throughput measurements for the report day. The Cisco router with the highest throughput is listed first. The Cisco router's throughput rank for the report day and the previous day (Previous Rank) is also given, along with the report day's percentage of discards. The throughput is shown in packets per second and bytes per second.

Cisco Routers With Highest Throughput
Sun Dec 19 1999

Router	Pkts/sec	Rank	Prev Rank	Bytes/sec	Pct Discards
boston	37741747.27	1	1	1213756965.47	1.86
tulsa	35379888.04	2	2	947004189.12	2.48
lincoln	33926123.71	3	3	443368212.95	3.73
miami	27480781.14	5	4	1182570806.32	3.73
houston	11812231.66	7	7	972845001.10	2.58
cotati	4015947.52	8	8	174540564.83	2.52

- The **Cisco Routers with Greatest Change in Throughput** table shows the Cisco routers with the largest change in rank from the preceding day to the report day. The Cisco router with the largest change in rank is listed first.

Cisco Routers With Greatest Change In Throughput
Sun Dec 19 1999

Router	Pkts/sec	Prev Pkts/sec	Change Rank	Rank	Prev Rank	Pct Discards
tulsa	35379888.04	35477323.17	2	2	2	2.48
miami	27480781.14	27384531.32	3	5	4	3.73
boston	37741747.27	37666101.63	4	1	1	1.86
houston	11812231.66	11834926.91	5	7	7	2.58
lincoln	33926123.71	33912132.44	6	3	3	3.73
cotati	4015947.52	4008552.08	8	8	8	2.52

- The **Cisco Routers with Highest Discard Percentage** table ranks the ten Cisco routers with the largest average percentage of discard scores for the report day.

Cisco Routers with Highest Discard Percentage
Sun Dec 19 1999

Device	Pct Pkts Discarded	Rank	Prior Rank	Pkts/sec	Bytes/sec
lincoln	3.73	1	1	33926123.71	443368212.95
miami	3.73	2	2	27480781.14	1182570806.32
houston	2.58	5	5	11812231.66	972845001.10
cotati	2.52	6	6	4015947.52	174540564.83
tulsa	2.48	7	7	35379888.04	947004189.12
boston	1.86	8	8	37741747.27	1213756965.47

- The **Cisco Routers with Greatest Change in Discard Percentage** table shows the Cisco routers with the greatest change in the average Discard Percentage score from the day preceding the report day to the report day. The Cisco router with the greatest change in Discard Percentage rank is listed first.

Cisco Routers With Greatest Change in Discard Percentage
Sun Dec 19 1999

Device	Pct Pkts Discarded	Previous Pct Pkts Discarded	Change Rank	Rank	Prior Rank	Pkts/sec	Bytes/sec
cotati	2.52	2.53	1	6	6	4015947.52	174540564.83
lincoln	3.73	3.74	2	1	1	33926123.71	443368212.95
boston	1.86	1.85	3	8	8	37741747.27	1213756965.47
miami	3.73	3.73	5	2	2	27480781.14	1182570806.32
houston	2.58	2.58	7	5	5	11812231.66	972845001.10
tulsa	2.48	2.48	8	7	7	35379888.04	947004189.12

Service Level Management Report

The Service Level Management report lets CIOs, CFOs, network managers, users and customers know how their Cisco routers are performing against contracted service levels. It reports the ten routers with lowest availability and the ten routers with the highest network response times. It also provides drill-down reports for daily

availability over the baseline period for the routers with lowest availability and daily and hourly response times for the routers with the highest network response time.

The Service Level Management Report comprises the following components:

- The **Availability** table lists (Cisco Routers with Lowest Availability for the Day) the Cisco routers with the lowest availability for yesterday. For more information about how availability is calculated, see “Daily Availability Drill-Down” on page 3-42.

  Availability
  Routers With Lowest Availability
  Sun Dec 19 1999

Device	Availability
miami	100.00
tulsa	100.00
cotati	100.00
lincoln	100.00
boston	100.00
houston	100.00

  Double-click on a Cisco router to update the content of the following components with performance data for the selected router:

- The **Daily Availability** line graph shows the average availability and service level of the selected Cisco router for each sample over the report period (42 days). By default, the service level is set at 99.5%.

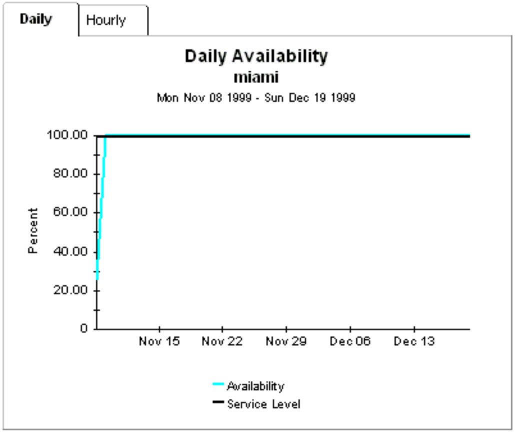

- The **Hourly Availability** line graph shows the average hourly availability and service level for the selected Cisco router for each sample for the preceding day. By default, the service level is set at 99.5%.

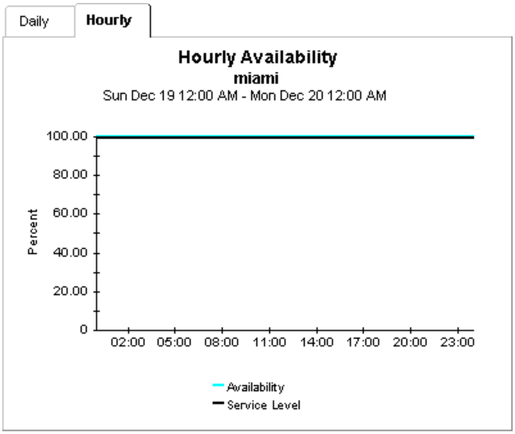

- The **Response Time** table lists the Cisco routers with the highest (average) response times for the day. For more information about how response time is calculated, see "Response Time" on page 3-45.

Response Time

Routers With Highest Response Time

Sun Dec 19 1999

Device	Response Time
houston	558.43
miami	500.00
boston	479.52
cotati	437.75
tulsa	362.31
lincoln	272.95

Double-click on a Cisco router to update the content of the following components with performance data for the selected router:

- The **Daily Response Time** line graph shows the average response time for the selected Cisco router for each day of the baseline period.

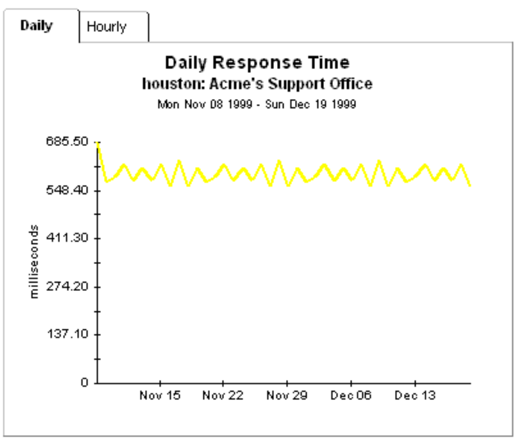

- The **Hourly Response Time** line graph shows the response time for each sample for the report day for the selected Cisco router.

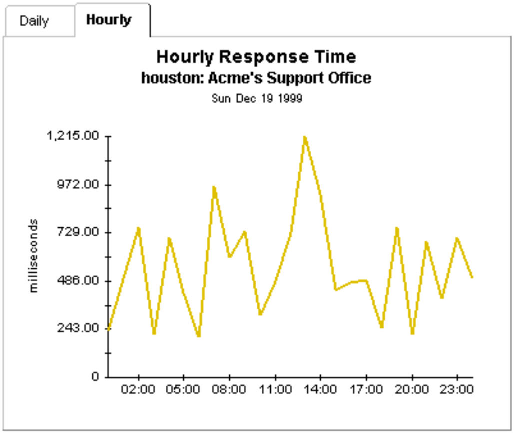

Data Source

Source data for the Cisco Router ReportPack is collected using SNMP from the Cisco 10 MIB.

Table 8-1 lists the variables that are polled by SNMP from the Cisco 10 MIB.

Table 8-1: Cisco Router Variables

Variable	**OID String**
bufferSmSize	1.3.6.1.4.1.9.2.1.14
Contains the size of small buffers.	
bufferSmTotal	1.3.6.1.4.1.9.2.1.15

(1 of 5)

Table 8-1: Cisco Router Variables

Variable	**OID String**
Contains the total number of small buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferSmFree	1.3.6.1.4.1.9.2.1.16
Contains the number of free small buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferSmHit	1.3.6.1.4.1.9.2.1.18
Contains the number of small buffer hits. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferSmMiss	1.3.6.1.4.1.9.2.1.19
Contains the number of small buffer misses. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferMdSize	1.3.6.1.4.1.9.2.1.22
Contains the size of medium buffers.	
bufferMdTotal	1.3.6.1.4.1.9.2.1.23
Contains the total number of medium buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferMdFree	1.3.6.1.4.1.9.2.1.24
Contains the number of free medium buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferMdHit	1.3.6.1.4.1.9.2.1.26
Contains the number of medium buffer hits. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	

(2 of 5)

Table 8-1: Cisco Router Variables

Variable	OID String
buffermdMiss	1.3.6.1.4.1.9.2.1.27
Contains the number of medium buffer misses. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferBgSize	1.3.6.1.4.1.9.2.1.30
Contains the size of big buffers.	
bufferBgTotal	1.3.6.1.4.1.9.2.1.31
Contains the total number of big buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferBgFree	1.3.6.1.4.1.9.2.1.32
Contains the number of free big buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferBgHit	1.3.6.1.4.1.9.2.1.34
Contains the number of big buffer hits. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferBgMiss	1.3.6.1.4.1.9.2.1.35
Contains the number of big buffer misses. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferLgSize	1.3.6.1.4.1.9.2.1.38
Contains the size of large buffers.	
bufferLgTotal	1.3.6.1.4.1.9.2.1.39
Contains the total number of large buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	

(3 of 5)

Table 8-1: Cisco Router Variables

Variable	OID String
bufferLgFree	1.3.6.1.4.1.9.2.1.40
Contains the number of free large buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferLgHit	1.3.6.1.4.1.9.2.1.42
Contains the number of large buffer hits. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferLgMiss	1.3.6.1.4.1.9.2.1.43
Contains the number of large buffer misses. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferFail	1.3.6.1.4.1.9.2.1.46
Count of the number of buffer allocation failures. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferNoMem	1.3.6.1.4.1.9.2.1.47
Count of the number of buffer create failures due to no free memory. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferHgSize	1.3.6.1.4.1.9.2.1.62
Contains the size of huge buffers.	
bufferHgTotal	1.3.6.1.4.1.9.2.1.63
Contains the total number of huge buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferHgFree	1.3.6.1.4.1.9.2.1.64

(4 of 5)

Table 8-1: Cisco Router Variables

Variable	OID String
Contains the number of free huge buffers. Changed from INTEGER to Gauge by DeskTalk Systems, Inc.	
bufferHgHit	1.3.6.1.4.1.9.2.1.66
Contains the number of huge buffer hits. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	
bufferHgMiss	1.3.6.1.4.1.9.2.1.67
Contains the number of huge buffer misses. Changed from INTEGER to Counter by DeskTalk Systems, Inc.	

(5 of 5)

Basic Metric Calculations

The performance metrics used in Cisco Router reports are based on the calculations listed in Table 8-2. However, some charts modify these calculations, so be sure to check the metrics given for each chart in each report in subsequent sections of this chapter.

Table 8-2: Cisco Router Metrics

Metric	Computed As ...
CPU Utilization	cpubusy5
The ratio of the amount of CPU time used (collected as centiseconds and multiplied by 100 to yield seconds) to the maximum amount of CPU time available (collected in centiseconds and multiplied by 100 to yield seconds) is multiplied by 100 to yield a percentage.	

Table 8-2: Cisco Router Metrics

Metric	Computed As ...
Memory Utilization	physical_mem_util = ((1 - (freememory1/ physicalram)) * 100);
The ratio of total memory (in bytes) used to total memory available (in bytes) is multiplied by 100 to yield a percentage.	
Buffer Utilization	memp1_util = (memp1_free/memp1_count)*100.0 (small) memp2_util = (memp2_free/memp2_count)*100.0 (medium) memp3_util = (memp3_free/memp3_count)*100.0 (large) memp4_util = (memp4_free/memp4_count)*100.0 (big) memp5_util = (memp5_free/memp5_count)*100.0 (huge)
The ratio of the number of free buffers to the total number of buffers times 100 to yield a percentage.	

Grade of Service Calculation

Table 8-3 lists how the metrics are weighted in Cisco Router Grade of Service (GOS) charts

Table 8-3: Cisco Router GOS Calculations

GOS Calculation	CPU Utilization	Memory Utilization	Buffer Utilization	Buffer Faults
Weighting	10%	40%	20%	30%
"Excellent" Score Range	0-19%	0-19%	0-19%	0-99
"Good" Score Range	20-39%	20-39%	20-39%	100-399
"Warning" Score Range	40-79%	40-79%	40-79%	400-799
"Critical" Score Range	80-100%	80-100%	80-100%	800 and over

For more information about the Grade of Service metric, see Table 2-2 on page 2-6. For more information about Grade of Service as it applies to each type of report, see "Analyzing TREND Reports" on page 3-1.

TREND

9 RMON ReportPack

The RMON ReportPack helps TREND users manage the various segments of their Ethernet networks, Token Ring networks, or both. It helps identify possible trouble spots and pinpoint what may have caused the problem or point out a situation that may become a problem.

Remote Monitoring (RMON) is an IEFT standard SNMP Management Information Base designed to provide network performance information remotely. Typically, it is implemented by hardware or software that promiscuously reads packets off of the network, scans their header information, and collects a large number of statistics about traffic crossing the network. These statistics include the number and size of packets traversing the network by protocol, source, and destination. The RMON agent provides this information using SNMP and, in some cases, proprietary protocols.

The ReportPack provides four QuickView reports. The reports help you identify networks that are overutilized, have a high number of collisions, or both. These factors can indicate that a segment is overutilized. A possible response to an overutilized segment is to create another segment. A low utilization level combined with a high percentage of collisions could indicate a bad Ethernet card on the segment, which would require replacing the card.

Understanding Remote Monitoring (RMON)

RMON tracks network traffic that involves device other than the host device. The monitoring can be local or remote.

Local Monitoring

Most network devices can record events that affect them. This is often done by an agent of some kind, typically, an SNMP agent. An SNMP agent on a router can track the actions and performance on that router. However, such an agent cannot typically record what is occurring on other devices. Additionally, while the same agent can track how many packets were received and sent by the router, it has no way to identify from which device the packets were received or to which device the packets were sent. Thus, local monitoring can tell you a great deal about the performance of a particular device but less about the health of a network segment.

Remote Monitoring

Unlike a local agent, an RMON agent is not typically concerned with the actions of its host device. Instead, an RMON device monitors a network segment, recording all of the traffic it sees. An RMON device is designed to track conversations between devices—who sent what to whom—and when and how those conversations took place. This provides more definitive information on how segment bandwidth is used and can highlight potential trouble spots.

Report Descriptions

The following sections describe each report in the RMON ReportPack in detail.

QuickView Reports

The QuickView reports provide daily and monthly information on individual Ethernet or Token Ring networks. The daily reports help you identify potentially critical short-term problems on your network. The monthly reports help you evaluate utilization trends for your networks.

Daily RMON Ethernet

The **Daily RMON Ethernet** report lists each Ethernet interface with its corresponding average and peak utilization and average and peak percent collisions for the preceding day. It comprises the following components:

- The **Daily RMON Ethernet Statistics** table provides a list of all Ethernet interfaces with their daily average and peak utilization and average and peak percentage collisions. The interfaces are sorted in descending order based on

daily peak utilization. Percent collisions is calculated by dividing the number of collisions by the number of good packets.

Daily RMON Ethernet Statistics
Sorted by Peak Utilization
Wed Oct 20 1999

RMON Probe	I/F	Avg Util	Peak Util	Avg Pct Collisions	Peak Pct Collisions
bluegrass	1_ALL	3.86	57.38	0.02	0.17
fms-east.desktalk.com	1_ALL	1.95	23.64	0.01	0.02
fp2000-east.desktalk.com	1_ALL	1.90	22.52	2.23	0.74
fp2000-east.desktalk.com	2_ALL	1.89	22.42	0.44	0.37
sanford	1_ALL	3.04	14.05	0.02	0.10
bluegrass	3_ALL	0.39	1.13	0.00	0.00
bluegrass	110_BI_ALL	0.39	1.13	0.00	0.00
bluegrass	110_DTE_ALL	0.21	0.80	0.00	0.00
bluegrass	21_ALL	0.21	0.80	0.00	0.00

Double-click on an interface in the table to update the content of the following components with performance data for the selected interface:

- The **RMON Ethernet Utilization** line graph shows the daily average, 95th percentile, and peak utilization for the past 30 days, if available, for the interface selected in the **Daily RMON Ethernet Statistics** table.

RMON Ethernet Utilization
bluegrass: 1_ALL
Tue Sep 21 1999 - Wed Oct 20 1999

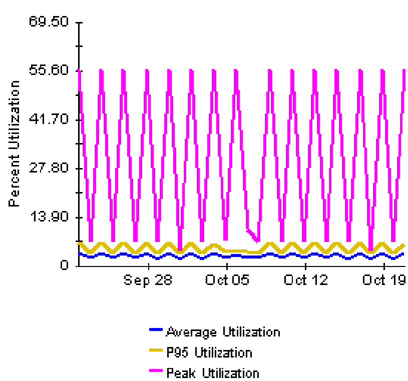

- The **RMON Ethernet Collisions** line graph shows the daily average and peak percent collisions for the past 30 days, if available, for the interface selected in the **Daily RMON Ethernet Statistics** table.

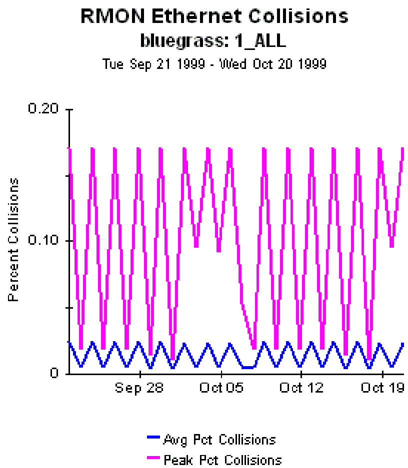

Monthly RMON Ethernet

The **Monthly RMON Ethernet** report lists each Ethernet interface with its corresponding average and peak utilization and average and peak percent collisions for the preceding month. It comprises the following components:

- The **Monthly RMON Ethernet Statistics** table provides a list of all Ethernet interfaces with their monthly average and peak utilization and average and peak percentage collisions. The interfaces are sorted in descending order based on monthly peak utilization. Percent collisions is calculated by dividing the number of collisions by the number of good packets.

Monthly RMON Ethernet Statistics
Sorted by Peak Utilization
Aug 1999

RMON Probe	I/F	Avg Pct Util	Peak Pct Util	Avg Pct Collisions	Peak Pct Collisions
bluegrass	1_ALL	3.00	56.15	0.01	0.17
sanford	1_ALL	3.22	27.19	0.02	0.07
fms-east.desktalk.com	1_ALL	1.78	22.67	0.01	0.02
fp2000-east.desktalk.com	1_ALL	1.73	21.57	2.15	0.74
fp2000-east.desktalk.com	2_ALL	1.70	21.48	0.47	0.37
bluegrass	3_ALL	0.32	1.08	0.00	0.00
bluegrass	110_BI_ALL	0.32	1.08	0.00	0.00
bluegrass	110_DTE_ALL	0.17	0.78	0.00	0.00
bluegrass	21_ALL	0.17	0.78	0.00	0.00

Double-click on an interface in the table to update the content of the following components with performance data for the selected interface:

- The **RMON Ethernet Utilization** bar chart shows the monthly average, 95th percentile, and peak utilization for the past 12 months, if available, for the interface selected in the **Monthly RMON Ethernet Statistics** table.

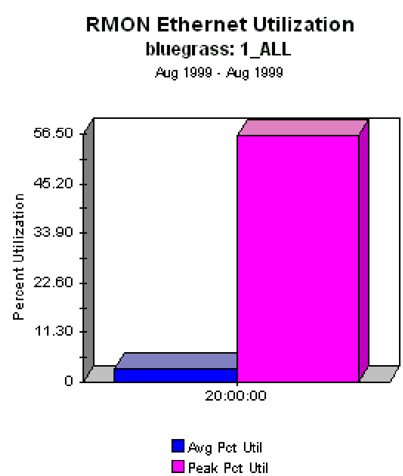

- The **RMON Ethernet Collisions** bar chart shows the monthly average and peak percent collisions for the past 12 months, if available, for the interface selected in the **Monthly RMON Ethernet Statistics** table.

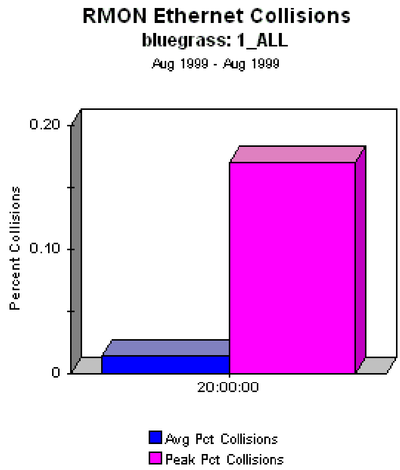

Daily RMON Token Ring

The **Daily RMON Token Ring** comprises the following components:

- The **Daily RMON Token Ring Statistics** table provides a list of all Token Ring interfaces with their daily average and peak utilization, average beacons

per second, and total beacons. The interfaces are sorted in descending order based on daily peak utilization.

Daily RMON Token Ring Statistics
Sorted by Peak Utilization
Wed Oct 13 1999 - Wed Oct 13 1999

RMON Probe	I/F	Speed	Avg Pct Util	Peak Pct Util	Avg Beacons / Sec	Total Beac
bluegrass	3_ALL-64000	64000.0	0.56	1.87	0.00	0.00
bluegrass	110_BI_ALL-64000	64000.0	0.56	1.86	0.00	0.00
bluegrass	21_ALL-64000	64000.0	0.30	1.32	0.00	0.00
bluegrass	110_DTE_ALL-64000	64000.0	0.30	1.32	0.00	0.00
bluegrass	22_ALL-64000	64000.0	0.26	0.86	0.00	0.00
bluegrass	110_DCE_ALL-64000	64000.0	0.26	0.85	0.00	0.00
bluegrass	1_ALL-10000000	1.0E7	0.04	0.62	0.03	2509.67
fms-east.desktalk.com	1_ALL-10000000	1.0E7	0.02	0.25	0.01	457.33

Double-click on an interface in the table to update the content of the following components with performance data for the selected interface:

- The **RMON Token Ring Utilization** line graph shows the daily average, 95th percentile, and peak utilization for the past 30 days, if available, for the interface selected in the **Daily RMON Token Ring Statistics** table.

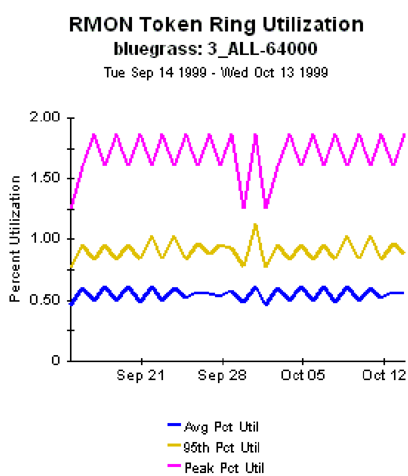

- The **RMON Token Ring Beacons** line graph shows the daily average and peak beacons per second for the past 30 days, if available, for the interface selected in the **Daily RMON Token Ring Statistics** table.

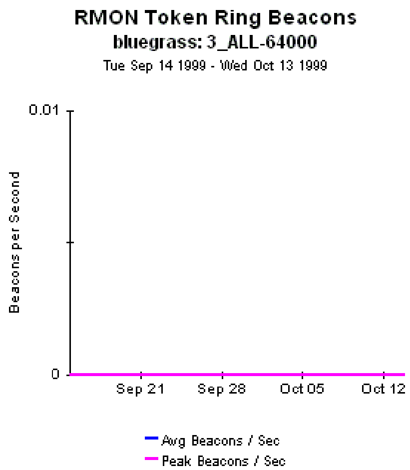

Monthly RMON Token Ring

The **Monthly RMON Token Ring** report comprises the following components:

- The **Monthly RMON Token Ring Statistics** table provides a list of all Token Ring interfaces with their monthly average and peak utilization, average

beacons per second, and total beacons. The interfaces are sorted in descending order based on monthly peak utilization.

Monthly RMON Token Ring Statistics
Sorted by Peak Utilization
Aug 1999

RMON Probe	I/F	Speed	Avg Pct Util	Peak Pct Util	Avg Beacons / Sec	Total Beacons
bluegrass	3_ALL-64000	64000.0	0.55	1.87	0.00	0.00
bluegrass	110_BI_ALL-64000	64000.0	0.55	1.86	0.00	0.00
bluegrass	21_ALL-64000	64000.0	0.30	1.35	0.00	0.00
bluegrass	110_DTE_ALL-64000	64000.0	0.29	1.35	0.00	0.00
bluegrass	22_ALL-64000	64000.0	0.26	1.00	0.00	0.00
bluegrass	110_DCE_ALL-64000	64000.0	0.26	1.00	0.00	0.00
bluegrass	1_ALL-10000000	1.0E7	0.03	0.62	0.02	37020.65
fms-east.desktalk.com	1_ALL-10000000	1.0E7	0.02	0.25	0.00	10046.19
fp2000-east.desktalk.com	1_ALL-10000000	1.0E7	0.02	0.24	1.23	2894206.53

Double-click on an interface in the table to update the content of the following components with performance data for the selected interface:

- The **RMON Token Ring Utilization** bar chart shows the monthly average, 95th percentile, and peak utilization for the past 12 months, if available, for the interface selected in the **Monthly RMON Token Ring Statistics** table.

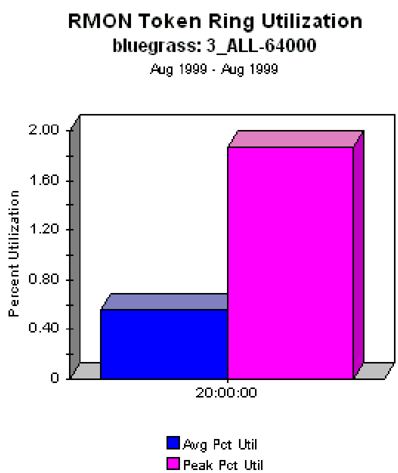

- The **RMON Token Ring Beacons** bar chart shows the monthly average and peak beacons per second for the past 12 months, if available, for the interface selected in the **Monthly Token Ring Ethernet Statistics** table.

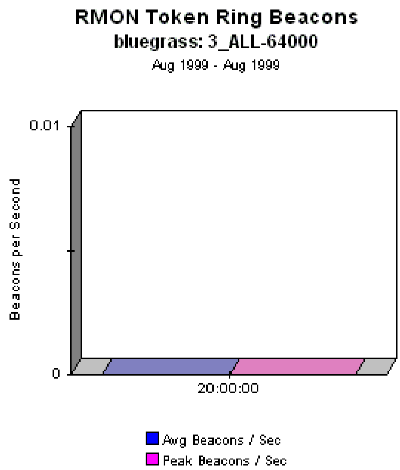

TREND

10 RMON2 ReportPack

The RMON2 ReportPack summarizes data collected from RMON2 probes, which provides the network manager with information about network performance.

The ReportPack provides four reports (Hot Spots, QuickView, Snapshot, and Top Ten) that allow you to identify which hosts utilize the majority of your network segment's bandwidth. Additionally, they highlight slow network response time by showing your segments with high utilization, error counts, or both. This information can help you reorganize your segments based on host traffic patterns.

Report Descriptions

The following sections describe each report in the RMON2 ReportPack in detail.

Hot Spots Report

This report lists all network interfaces with utilization greater than 20% and/or collisions that amount to greater than 15% of good packets for the preceding day. Associated charts provide hourly segment utilization, collisions, errors, and frame size distribution for the selected interface. For an in-depth discussion of the Hot Spots report, see "Hot Spots Report" on page 3-32.

The Hot Spots report comprises the following charts:

- The **Problem Summary for the Day** table provides a list of interfaces that meet or exceed the specified criteria, which is analyzed in related charts. The utilization specified is a daily average utilization. The percent collisions is calculated by dividing the total number of collisions for the day by the sum of the total number of dropped and transmitted packets.

Problem Summary for the Day

Segments with Average Utilization > 20 % or Collisions > 15 % of Good Packets

Mon Dec 13 1999

RMON Probe	I/F	Type	Speed	Utilization	Pct Collisions
bluegrass	3_ALL	frame-relay	64000.00	46.53	0.00
bluegrass	1_LHS_GEN	ethernet-csmacd	10000000.00	0.00	3774.53
fp2000-east.desktalk.com	1_ALL	ethernet-csmacd	10000000.00	1.73	149.87

Double-click on an interface to update the content of the following components with performance data for the interface:

- The **Segment Utilization** line graph shows the average and peak utilization by hour for the interface selected from the **Problem Summary for a Day** table.

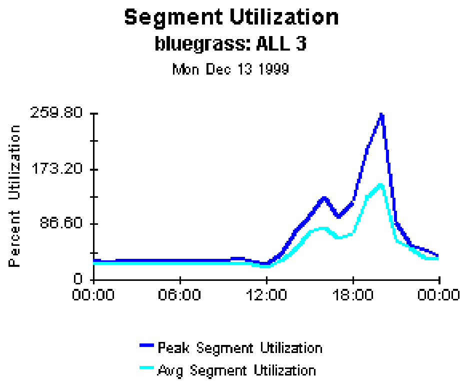

- The **Segment Collisions** line graph shows the total collisions, dropped events, and dropped frames by hour for the interface selected from the **Problem Summary for a Day** table.

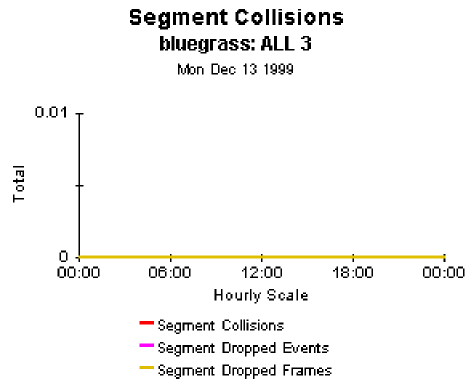

10 RMON2 ReportPack

- The **Segment Errors** line graph shows the total errors by type by hour for the interface selected from the **Problem Summary for a Day** table.

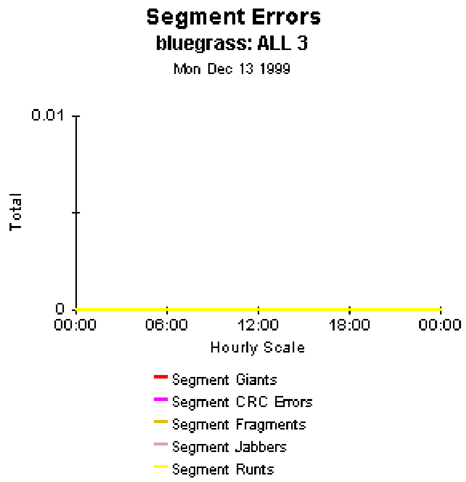

- The **Frame Size** stacked bar chart shows the total number of frames and the distribution of frame size by hour for the interface selected from the **Problem Summary for a Day** table.

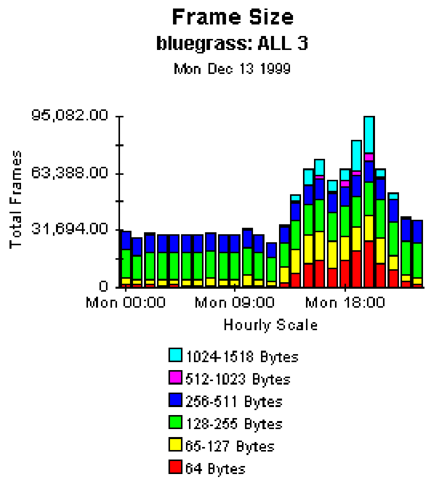

QuickView and Snapshot Reports

The QuickView report provides detailed information on individual interfaces in a single report. Specifically, it lets you determine if your interface has possible probe resource problems due to high collision rates, errors occurring on segments, or high segment utilization. The Snapshot report provides the same information as the Quick View report, however, it allows you to show detailed information about a particular interface. The selected interface may be different than the interfaces appearing in the QuickView report.

Daily QuickView Report

The **Daily QuickView** report lists each network interface with its corresponding average and peak utilization for the preceding day. The associated charts show the history of daily segment utilization, collisions, errors, and frame size distribution for the selected interface.

The **Daily QuickView** report comprises the following components:

- The **Daily Segment Utilization** table lists of all polled interfaces that you can select for further analysis in the associated charts. Both the daily average utilization and the daily peak utilization are specified. The interfaces are sorted in descending order based on the daily average utilization.

Daily Segment Utilization
Sorted by Average Utilization
Mon Dec 13 1999

RMON Probe	I/F	Type	Speed	Average Utilization	Peak Utilization
bluegrass	3_ALL	frame-relay	64000.00	46.53	259.64
fp2000-east.desktalk.com	1_ALL	ethernet-csmacd	10000000.00	1.73	15.33
bluegrass	1_LHS_GEN	ethernet-csmacd	10000000.00	0.00	0.00

Double-click on an interface in the chart to update the content of the following components with performance data for the selected interface:

- The **Daily Utilization** line graph shows the daily average and peak utilization for the past six weeks, if available, for the interface selected from the **Daily Segment Utilization** table.

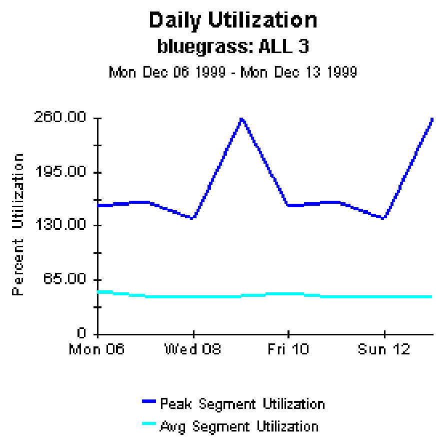

- The **Collisions, Dropped Frames and Events** line graph shows the daily total collisions, dropped events, and dropped frames for the past six weeks, if available, for the interface selected from the **Daily Segment Utilization** table.

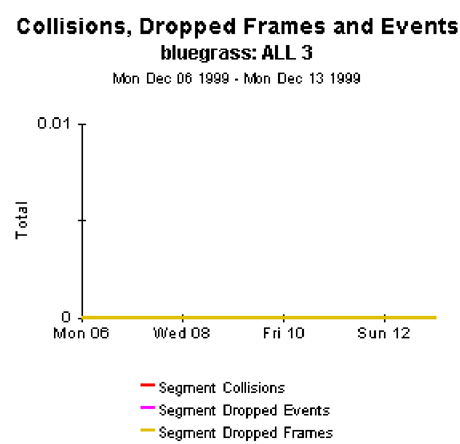

- The **Segment Errors** line graph shows the daily total errors by type for the past six weeks, if available, for the interface selected from the **Daily Segment Utilization** table.

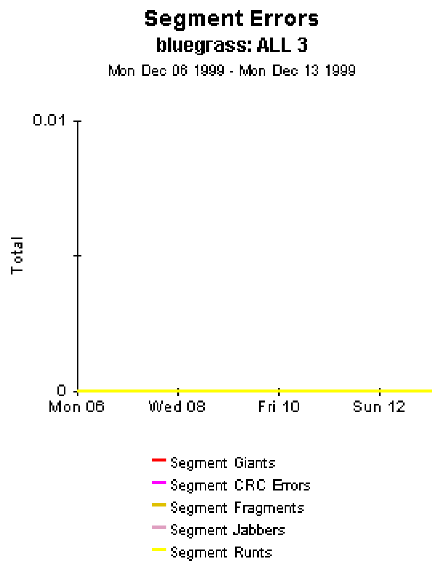

- The **Frame Size Distribution** stacked bar chart shows the daily total number of frames and the distribution of frame size for the past six weeks,

if available, for the interface selected from the **Daily Segment Utilization** table.

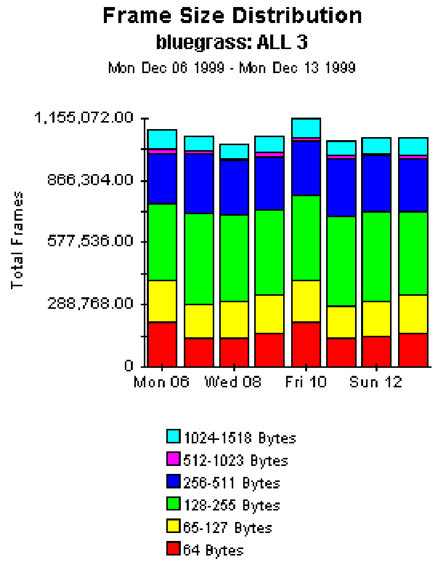

Hourly QuickView Report

The **Hourly QuickView** report lists each network interface with its corresponding average and peak utilization for the preceding day. Associated charts show the hourly breakdown of segment utilization, collisions, errors, and frame size distribution for the selected interface.

The **Hourly QuickView** report comprises the following components:

- The **Segment Utilization** table lists the polled interfaces that you can select for further analysis in the associated charts. Both the daily average utilization and the daily peak utilization are specified. The interfaces are sorted in descending order based on the daily average utilization.

Double-click on an interface in the chart to update the content of the following components with performance data for the selected interface:

- The **Hourly Utilization** line graph shows the average and peak utilization by hour for the interface selected from the **Segment Utilization** table.
- The **Collisions, Dropped Frames and Events** line graph shows the total collisions, dropped events, and dropped frames by hour for the interface selected from the **Segment Utilization** table.
- The **Segment Errors** line graph shows the total errors by type by hour for the interface selected from the **Segment Utilization** table.
- The **Frame Size Distribution** stacked bar chart shows the total number of frames and the distribution of frame size by hour for the interface selected from the **Segment Utilization** table.

Snapshot Reports

The RMON2 ReportPack provides the following Snapshot reports:

- **Daily Snapshot**
- **Hourly Snapshot**

They provide the same information as their corresponding QuickView reports; however, you can select the specific interfaces that you want the report to include. For more information about selecting an interface for a Snapshot report, see "Selecting Deferred Targets or Elements for a Snapshot Report" on page 1-27.

Top Ten Report

The Top Ten reports list the top ten conversation pairs, talkers, and listeners for the preceding day and preceding week. Specifically, it lets you determine if your interfaces are not optimally deployed, have a hardware problem, or overloaded by requests.

Daily Top Ten

The **Daily Top Ten** report lists the top-ten conversation pairs, talkers, and listeners for the preceding day.

The **Daily Top Ten** report comprises the following components:

- The **Top Conversation Pairs** table lists the source and destination IP address pairs that sent the most data on the preceding day along with the total megabytes sent.

Top Converstion Pairs
Mon Dec 13 1999

Source Address	Target Address	Sent Mbytes
134.70.65.84	134.70.32.251	45.13
134.70.32.24	134.70.150.17	25.08
134.70.150.17	134.70.32.24	15.32
134.70.32.5	209.219.153.173	14.91
134.70.32.5	134.70.65.206	14.50
134.70.32.5	134.70.180.100	14.50
134.70.32.5	134.70.65.106	14.42
134.70.65.206	134.70.32.5	14.18
134.70.65.106	134.70.32.5	14.13
209.219.153.173	134.70.32.5	13.94

- The **Top Ten Talkers** table lists the IP address and the total megabytes sent and received for the addresses that sent the most data to all destinations on the preceding day.

Top Ten Talkers
Mon Dec 13 1999

Src Address	Probe	IF Number	Network Layer	Sent Mbytes	Rcv Mbytes
134.70.32.5	bluegrass	3	ip[1]	72.49	72.02
134.70.65.84	bluegrass	3	ip[1]	43.48	3.01
134.70.150.17	bluegrass	3	ip[1]	26.09	29.32
134.70.32.24	bluegrass	3	ip[1]	25.44	17.43
134.70.65.106	bluegrass	3	ip[1]	16.01	15.85
209.219.153.173	bluegrass	3	ip[1]	15.47	15.91
134.70.180.100	bluegrass	3	ip[1]	15.02	15.46
134.70.65.206	bluegrass	3	ip[1]	14.89	15.26
134.70.180.113	bluegrass	3	ip[1]	13.93	14.32
134.70.32.1	bluegrass	3	ip[1]	8.77	7.66

- The **Top Ten Listeners** table lists the IP address and the total megabytes sent and received for the addresses that received the most data from all destinations on the preceding day.

Top Ten Listeners
Mon Dec 13 1999

Src Address	Probe	IF Number	Network Layer	Rcv Mbytes	Sent Mbytes
134.70.32.5	bluegrass	3	ip[1]	72.02	72.49
134.70.32.251	bluegrass	3	ip[1]	52.70	4.47
134.70.150.17	bluegrass	3	ip[1]	29.32	26.09
134.70.32.24	bluegrass	3	ip[1]	17.43	25.44
209.219.153.173	bluegrass	3	ip[1]	15.91	15.47
134.70.65.106	bluegrass	3	ip[1]	15.85	16.01
134.70.180.100	bluegrass	3	ip[1]	15.46	15.02
134.70.65.206	bluegrass	3	ip[1]	15.26	14.89
134.70.180.113	bluegrass	3	ip[1]	14.32	13.93
134.70.32.56	bluegrass	3	ip[1]	12.35	0.82

Weekly Top Ten

The **Weekly Top Ten** report lists the top-ten conversation pairs, talkers, and listeners for the preceding week.

The **Weekly Top Ten** report comprises the following components:

- The **Top Conversation Pairs** table lists the source and destination IP address pairs that sent the most data on the preceding week along with the total megabytes sent.
- The **Top Ten Talkers** table lists the IP address and the total megabytes sent and received for the addresses that sent the most data to all destinations on the preceding week.
- The **Top Ten Listeners** table lists the IP address and the total megabytes sent and received for the addresses that received the most data from all destinations on the preceding week.

11 TREND Database ReportPack

The TREND Database ReportPack reports historical and predicted usage levels of the TREND database and transaction log. It provides a tool for the TREND database administrator that shows trends in the database use, and also a means to predict when to increase the database capacity of the system.

The TREND Database ReportPack answers questions including:

- Do I need to increase the size of the database for better utilization?
- Do I have enough space to load new applications or ReportPacks?

The report consists of a combination of charts and graphs. To produce the report, the TREND Database ReportPack uses data from the TREND processes in their normal mode of operation. There is no data collection for this report, in the normal sense.

Report Descriptions

The following section describes the report in the TREND Database ReportPack in detail.

Database Usage Summary Report

The Database Usage Summary report enables the TREND administrator to view the historic utilization of the TREND database and transaction log as well as view forecasts for 30, 60, and 90 days into the future. The TREND administrator can use this information to better utilize disk space while TREND continues to collect data.

The forecasts use the data from the baseline period, which by default is the previous 42 days (or 6 weeks), to predict the usage level for the next, 30, 60, or 90 days.

The Database Usage Summary report comprises the following components:

- The **Database Utilization** line graph shows the percentage of the daily average and maximum database utilization for the last 90 days.

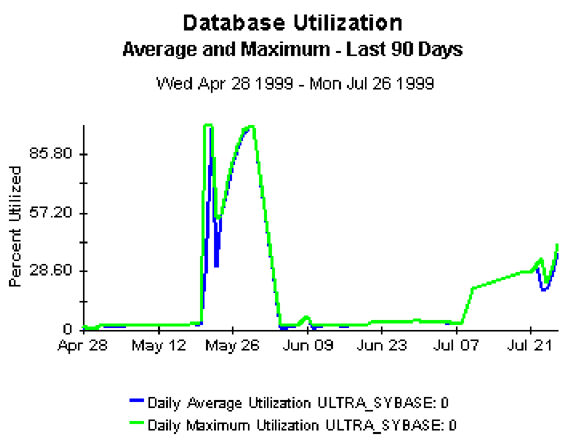

- The **Database Average Utilization** bar chart shows the average percentage of the database used for the baseline period and the projected percent usage of the database for the 30-, 60-, and 90-day forecasts.

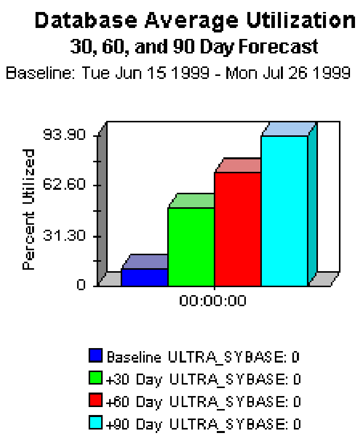

- The **Transaction Log Utilization** line graph shows the percentage of the daily average and maximum transaction log utilization for the last 90 days.

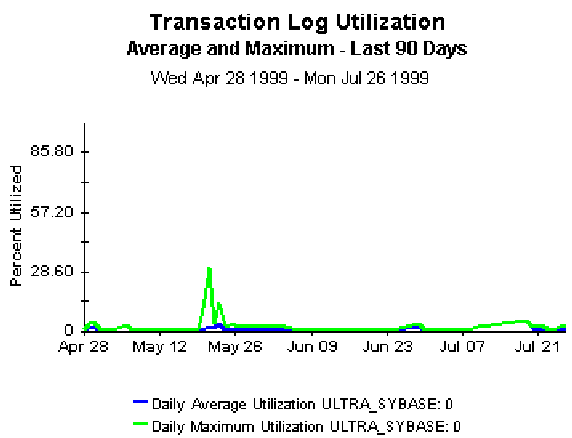

- The **Transaction Log Max Utilization** bar chart shows the average percentage of the transaction log used for the baseline period and the projected percent usage of the transaction log for the 30-, 60-, and 90-day forecasts.

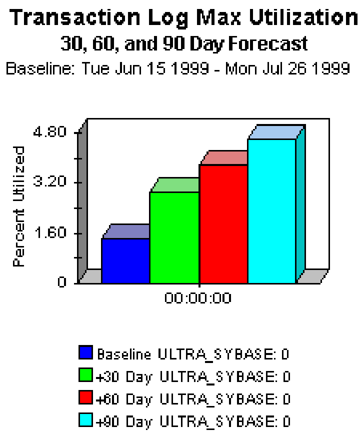

- The **Database Usage and Size** line graph shows the number of kilobytes for the daily average and maximum database usage for the last 90 days. It also shows the database size in kilobytes.

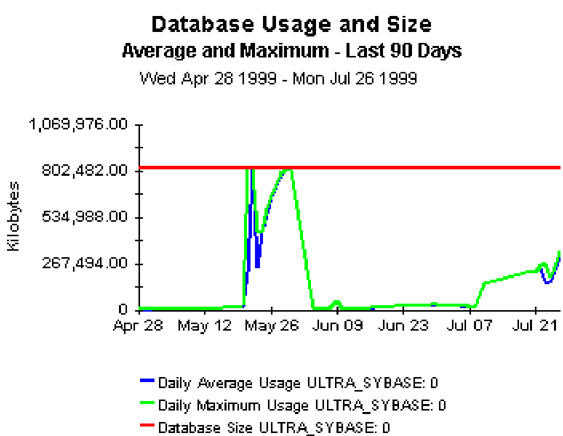

- The **Database Average Usage** bar chart shows the average number of kilobytes the database used for the baseline period and the number of kilobytes the database will use for the 30-, 60-, and 90-day forecasts.

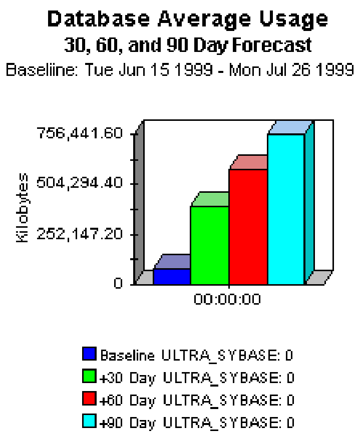

- The **Transaction Log Usage and Size** line graph shows the number of kilobytes for the daily average and maximum transaction log utilization for the last 90 days. It also shows the transaction log size in kilobytes.

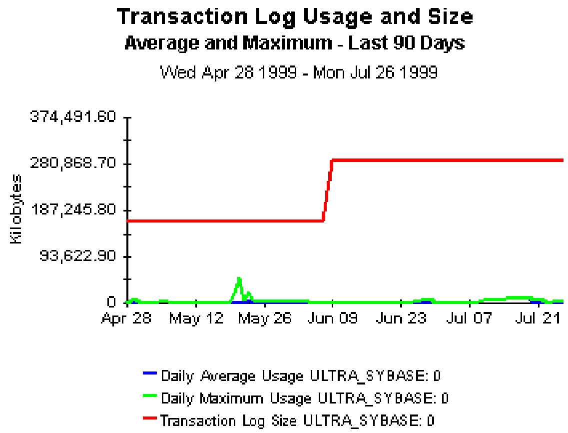

- The **Transaction Log Usage** bar chart shows the average number of kilobytes of the transaction log used for the baseline period and the number of kilobytes the transaction log will use for 30-, 60-, and 90-day forecasts.

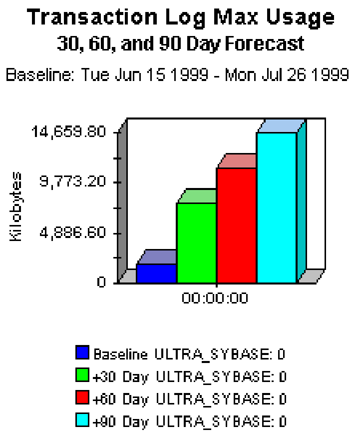

Even though the Utilization components look very similar to the Usage components, they actually show different information. The Utilization components change when the database or transaction log sizes change, but the Usage components stay the same since they show the actual use.

Data Source

The Usage Summary report uses the daily and forecast database statistics generated from TRENDit and TRENDsum. TRENDit populates the rate data table, Rsi_dbstats, from the data inserted into the raw table, dsi_dbstats, by the TREND database client applications. TRENDsum populates the SDdsi_database_util and SD42SDdsi_database_util_forecast tables that contain the daily and forecast data.

Source data for the TREND Database ReportPack is collected from the following tables:

- SDdsi_database_util table
- SD42SDdsi_database_util_forecast table

SDdsi_database_util Table

For the SDdsi_database_util table, TRENDsum calculates the daily average and maximum database size, database used, percentage of database used, transaction log size, transaction log used, and percentage of transaction log used. Table 11-1 lists the variables.

Table 11-1: TREND Database Daily Average and Maximum Variables

Variable	**Description**
AVGdbsizekb	The average size of the database in kilobytes, over each day.
MAXdbsizekb	The maximum size of the database in kilobytes, over each day.
AVGpct_dbused	The average percentage of the database that is used, over each day.
MAXpct_dbused	The maximum percentage of the database that is used, over each day.
AVGkb_dbused	The average number of kilobytes the database used, over each day.
MAXkb_dbused	The maximum number of kilobytes the database used, over each day.
AVGlogsizekb	The average size of the transaction log in kilobytes, over each day.

(1 of 2)

Table 11-1: TREND Database Daily Average and Maximum Variables

Variable	Description
MAXlogsizekb	The maximum size of the transaction log in kilobytes, over each day.
AVGpct_logused	The average percentage of the transaction log that is used, over each day.
MAXpct_logused	The maximum percentage of the transaction log that is used, over each day.
AVGkb_logused	The average number of kilobytes the transaction log used, over each day.
MAXkb_logused	The maximum number of kilobytes the transaction log used, over each day.

(2 of 2)

SD42SDdsi_database_util_forecast Table

For the SDdsi_database_util table, TRENDsum calculates the average, which is the baseline, and the 30-, 60-, 90-day forecast for database used and transaction log used. It also calculates the days to reach the 80% threshold for the percentage of database used and percentage of the transaction log used. Table 11-2 lists the variables.

Table 11-2: TREND Database Baseline and Forecast Variables

Variable	Description
AVGAVGkb_dbused	The average of daily averages for the number of kilobytes used by the database over the baseline period.

(1 of 3)

Table 11-2: TREND Database Baseline and Forecast Variables

Variable	**Description**
AVGAVGpct_dbused	The average of daily averages for the percentage of the database used over the baseline period.
AVGMAXkb_logused	The average of the daily maximum number of kilobytes used by the transaction log over the baseline period.
AVGMAXpct_logused	The average of the daily maximum percentage of the transaction log used over the baseline period.
DTTAVGpct_dbused	The projected number of days to reach the threshold of 80% of the daily average database usage. The ReportPack calculates but does not display this variable at this time.
DTTMAXpct_logused	The projected number of days to reach the threshold of 80% of the daily maximum transaction log usage. The ReportPack calculates but does not display this variable at this time.
F30AVGkb_dbused	The 30-day projection for the average number of kilobytes the database will use.
F30AVGpct_dbused	The 30-day projection for the average percentage of database use.
F30MAXkb_logused	The 30-day projection for the maximum number of kilobytes the transaction log will use.
F30MAXpct_logused	The 30-day projection for the maximum percentage of transaction log use.
F60AVGkb_dbused	The 60-day projection for the average number of kilobytes the database will use.

(2 of 3)

Table 11-2: TREND Database Baseline and Forecast Variables

Variable	**Description**
F60AVGpct_dbused	The 60-day projection for the average percentage of database use.
F60MAXkb_logused	The 60-day projection for the maximum number of kilobytes the transaction log will use.
F60MAXpct_logused	The 60-day projection for the maximum percentage of transaction log use.
F90AVGkb_dbused	The 90-day projection for the average number of kilobytes the database will use.
F90AVGpct_dbused	The 90-day projection for the average percentage of database use.
F90MAXkb_logused	The 90-day projection for the maximum number of kilobytes the transaction log will use.
F90MAXpct_logused	The 90-day projection for the maximum percentage of transaction log use.

(3 of 3)

Basic Metric Calculations

The performance metrics used in TREND Database reports are based on the elements from the dsi_dbstats table listed in Table 11-3. The table shows equivalent calculations for some of the elements.

Table 11-3: TREND Database Metrics

Metric	Computed As ...	How Computed
Database Utilization	prused	The percentage of database usage from dsi_dbstats, which is equivalent to the ratio of total bytes used to the size of the database, multiplied by 100.
Transaction Log Utilization	prlogused	The percentage of transaction log usage from dsi_dbstats, which is equivalent to the ratio of total bytes used to the size of the transaction log, multiplied by 100.
Database Usage	kbused	The total number of bytes used by the database.
Database Size	dbsizekb	The total number of bytes allocated for the database.
Transaction Log Usage	kblogused	The total number of bytes used by the transaction log.
Transaction Log Size	logsizekb	The total number of bytes allocated for the transaction log.

TREND

12 System ReportPack

The System ReportPack reports statistics about CPU utilization, paging and swapping, file system usage, virtual and physical memory, cache hits and misses, and grade of service specifically for the Empire SystemEDGE agent.

The System ReportPack answers questions including:

- Do I need to add more memory to my system?
- Am I running out of filesystem space?
- Is my system running low on resources?

Report Descriptions

The following sections describe each report in the System ReportPack in detail.

Executive Summary Report

The Executive Summary report provides high-level graphs aggregating key metrics for all systems. It allows executives to review the daily volume summary and, if desired, look into more detailed information such as grade of service, hourly volume compared to baseline, hourly grade of service, and total exception count.

The Executive Summary report comprises the following components:

- The **Daily Network Volume** bar chart shows the total volume for all systems over each day of the baseline period.

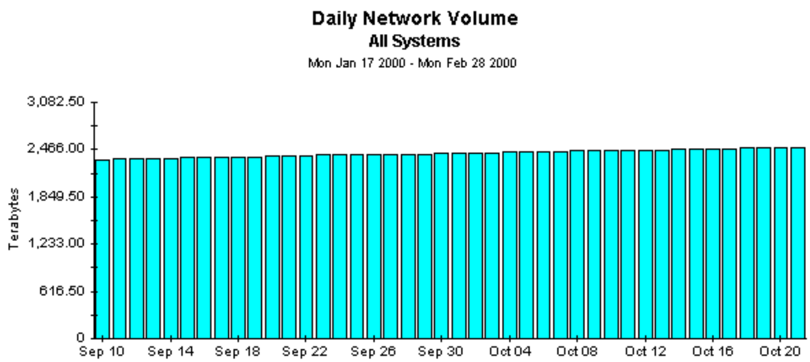

- The **Hourly Network Volume vs Baseline** bar chart shows the average volume for all systems for each hour of the report day compared to the baseline volume of all systems.

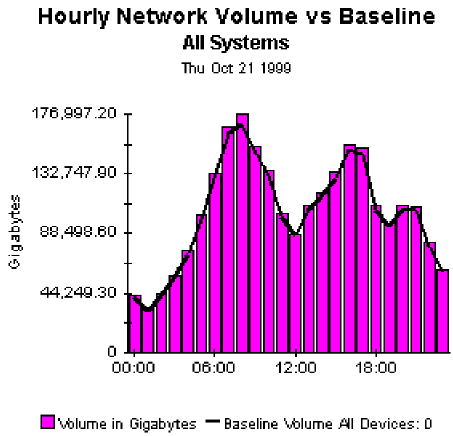

- The **Total Exception Counts** stacked bar chart shows the number of times a system exceeded a queue length or CPU, memory, or swap utilization threshold during each hour of the report day.

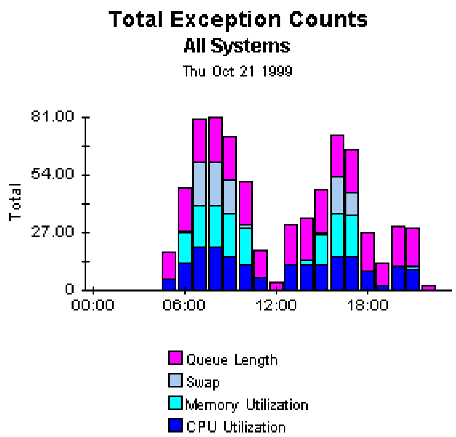

- The **Hourly Grade of Service** stacked bar chart provides an overview of system performance by supplying weighted values for queue length, and swap,

memory, and CPU utilization for each hour of the report day. Table 12-7 describes the GOS weighting factors used in this chart.

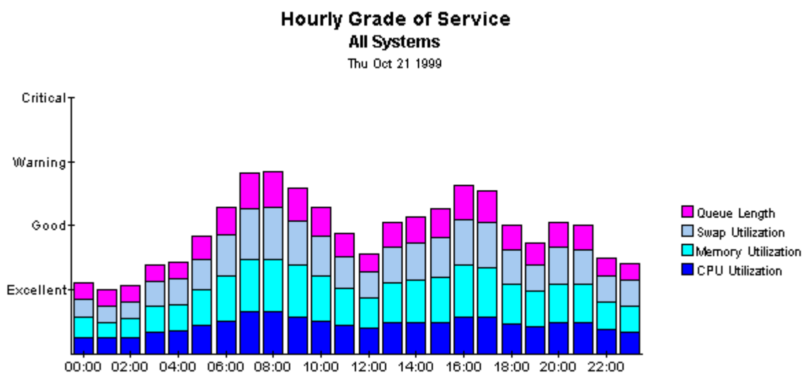

Forecast--Filesystems Report

The Forecast--Filesystems report lists a summary of all filesystems that are within 90 days of reaching a disk space utilization threshold. Those filesystems closest to the 90-day threshold appear in the list first. The calculation for the threshold is from the baseline period, which is the previous 42 days. The generation of the forecast data uses the 95th-percentile values from the daily exception table.

The report arms the network manager with the information required to assess performance and take preventive action. Each element in the summary hot links to multiple graphs that show CPU, memory, and swap utilization; grade of service; and page faults for trend analysis over the baseline period.

Note that there are usually multiple filesystems for each computer system; each filesystem references allocated disk space on the computer.

The Forecast--Filesystems report comprises the following components:

- The **Estimated Days to Threshold (DTT)** summary table lists the filesystems that are within 90 days of 100 percent utilization. It shows the days to threshold

for filesystem utilization, current filesystem utilization, and projected utilization in 90 days for each filesystem.

Estimated Days to Threshold (DTT)
Filesystems Within 90 Days Of 90% Utilization

Device	Filesystem	DTT Filesystem Util	Current Filesystem Utilization	Projected Util in 90 Days
wasp	C:	0.00	98.00	98.00
yorktown	/export/home	0.00	95.00	95.00

Double-click on a filesystem in the summary table to view the following components of performance information for the selection:

- The **Grade of Service** stacked bar chart shows weighted values for queue length, and CPU, memory, and swap utilization for each day of the baseline period for the selected filesystem. See Table 12-7 for a description of the GOS weighting factors.

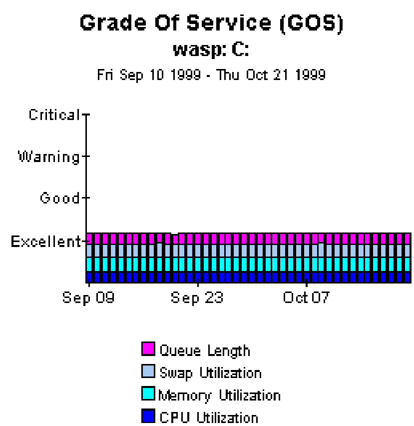

- The **Filesystem Utilization** line graph shows the average filesystem utilization for the selected filesystem for each day of the baseline period.

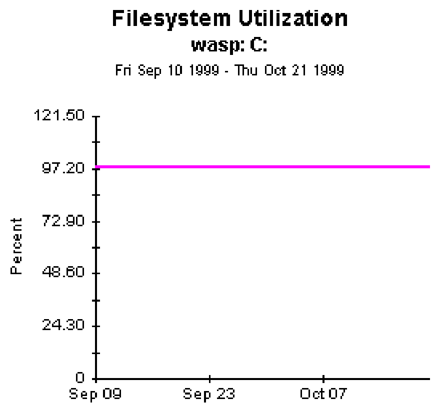

- The **CPU Utilization** line graph shows the average CPU utilization for the selected filesystem for each day of the baseline period.

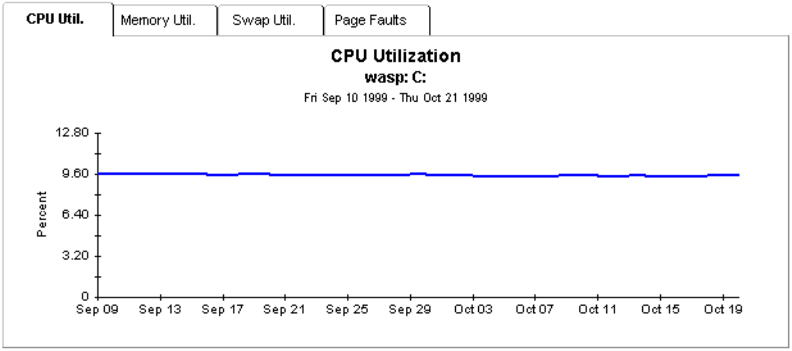

- The **Memory Utilization** line graph shows the average memory utilization for the selected filesystem for each day of the baseline period.

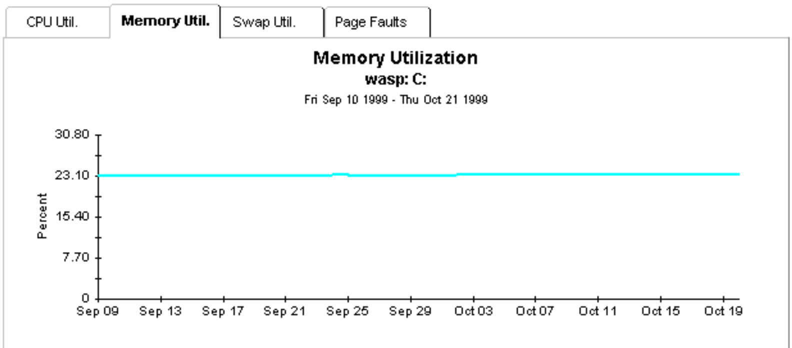

- The **Swap Utilization** line graph shows the average swap utilization for the selected filesystem for each day of the baseline period.

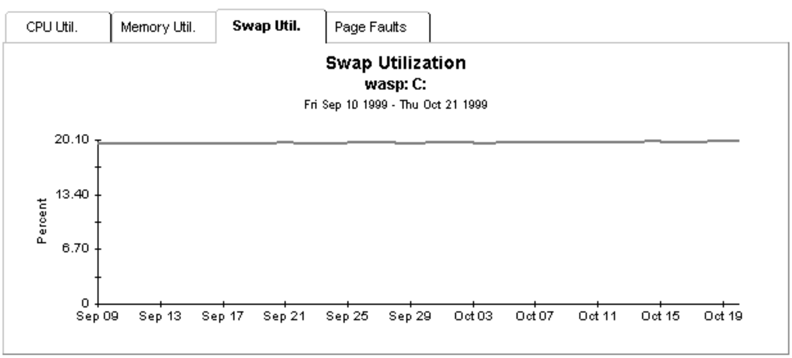

- The **Page Faults** bar chart shows the number of page faults for the selected filesystem for each day of the baseline period.

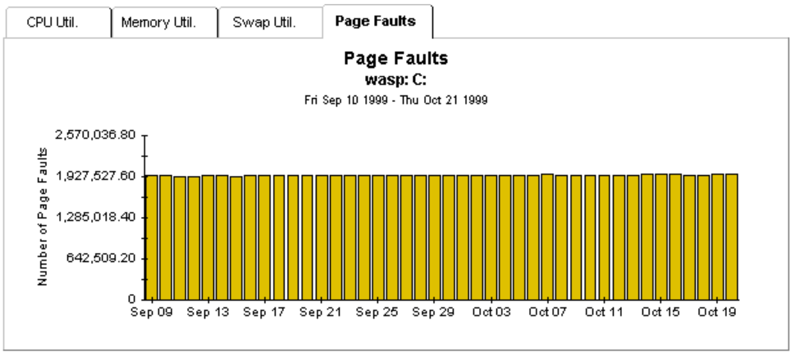

Forecast--Memory Report

The Forecast--Memory report lists all systems within 90 days of a memory utilization threshold. Those systems closest to a threshold appear in the list first. The report arms the network manager with the information required to assess performance and take preventive action. Each element in the summary hot-links to multiple graphs that show CPU, memory, and swap utilization; grade of service; total swaps; and page faults for trend analysis of the baseline performance period.

Note: The charts for the Forecast--Memory report are identical to the charts for the Forecast--Filesystems report. For chart illustrations, see "Forecast--Filesystems Report" on page 12-4.

The Forecast--Memory report comprises the following components:

- The **Estimated Days to Threshold (DTT)** summary table lists the systems that are within 90 days of a 100 percent memory utilization threshold. It shows the number of days to threshold for memory utilization, current memory

utilization, and projected memory utilization in 90 days for each computer system.

Double-click on a computer system in the summary table to view the following components of performance information for the selection:

- The **Grade of Service** stacked bar chart shows weighted values for CPU, memory, and swap utilization, and queue length for each day of the baseline period for the selected system. See Table 12-7 for a description of the GOS weighting factors.
- The **Memory Utilization** line graph shows the average memory utilization for the selected system for each day of the baseline period.
- The **CPU Utilization** line graph shows the average CPU utilization for the selected system for each day of the baseline period.
- The **Swap Utilization** line graph shows the average swap utilization for the selected system for each day of the baseline period.
- The **Total Swaps** bar chart shows the number of swaps for the selected system for each day of the baseline period.
- The **Page Faults** bar chart shows the number of page faults for the selected system for each day of the baseline period.

QuickView Report

The QuickView report provides the network management staff with detailed information on key metrics for individual systems with the highest GOS scores. It identifies the ten systems with the highest (worst) grade of service scores. Drilldowns show CPU utilization, memory utilization, swap utilization, page faults, and total swaps as well as number of processes running at the hourly level.

The QuickView report comprises the following components:

- The **Grade of Service** stacked bar chart shows the GOS scores for the selected systems. The chart shows weighted values for CPU, memory, and swap

utilization, and queue length for the report day. See Table 12-7 for a description of the GOS calculation formulas.

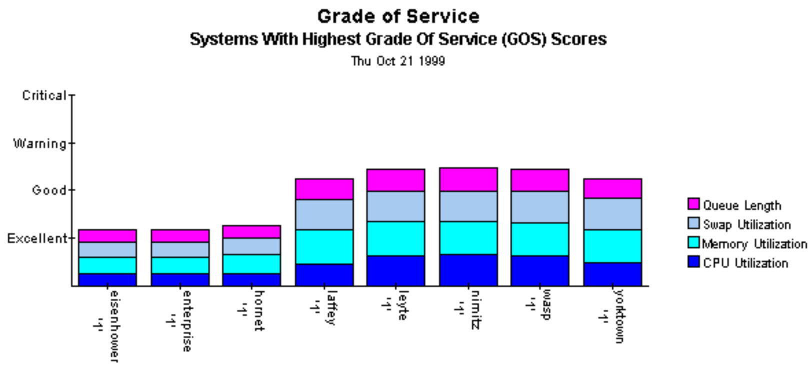

Double click on a system to update the following components with performance data for the selection:

- The **Number of Processes** line graph shows the number of processes for each sample period of the report day for the selected system.

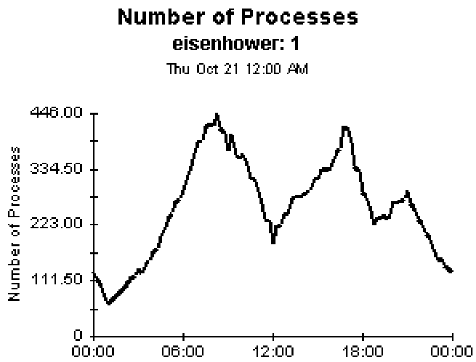

- The **CPU Utilization** line graph shows the CPU utilization for the selected system.

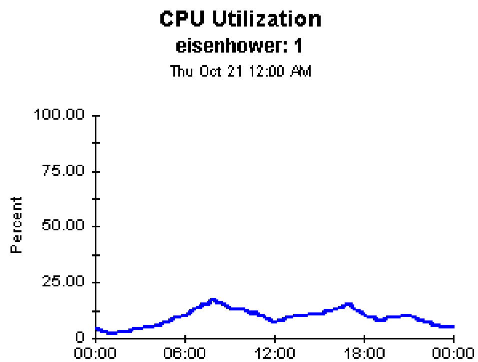

- The **Memory Utilization** line graph shows the memory utilization for the selected system.

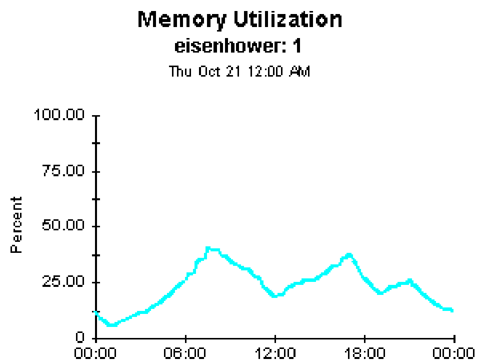

- The **Page Faults** bar chart shows the number of page faults for the selected system over each hour of the report day.

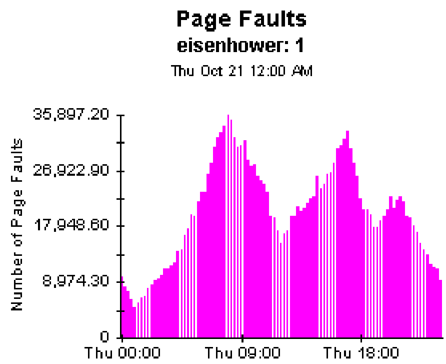

- The **Swap Utilization** line graph shows the swap utilization for the selected system.

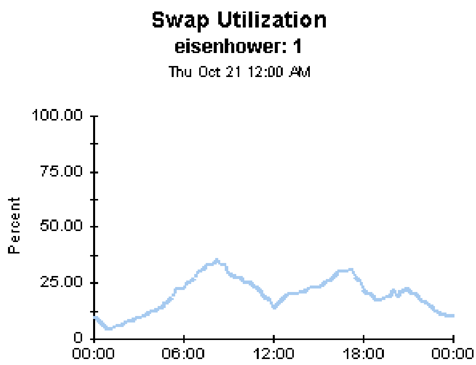

- The **Total Swaps** bar chart shows the number of swaps for the selected system over each hour of the report day.

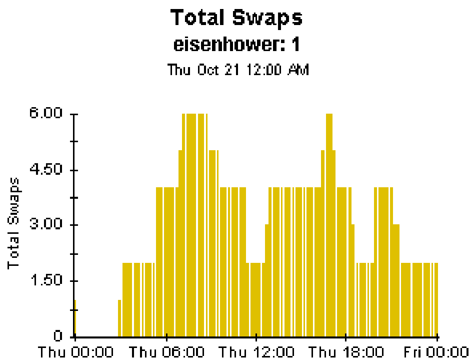

Snapshot Report

The Snapshot report shows the same information as the QuickView report for a system selected from a list at the time the report is invoked. For more information about selecting a System for a Snapshot report, see “Selecting Deferred Targets or Elements for a Snapshot Report” on page 1-27.

Note: The charts for the Snapshot report are identical to the charts for the QuickView report. For chart illustrations, see “QuickView Report” on page 12-9**.**

The Snapshot report comprises the following components:

- The Snapshot summary table shows the time period, device CPU utilization, memory utilization, total page faults, and swaps.

  Double click on a system to update the following components with performance data for the selection:

- The **Number of Processes** line graph shows the number of processes for each sample period of the report day for the selected system.
- The **CPU Utilization** line graph shows the CPU utilization for the selected system.
- The **Memory Utilization** line graph shows the memory utilization for the selected system.
- The **Page Faults** bar chart shows the number of page faults for the selected system over each hour of the report day.
- The **Swap Utilization** line graph shows the swap utilization for the selected system.
- The **Total Swaps** bar chart shows the number of swaps for the selected system over each hour of the report day.

Near Real Time QuickView and Snapshot Reports

The Near Real Time--QuickView report provides up-to-date system performance statistics up to the last SNMP poll. The report does not rely on nightly summaries; therefore, it provides instant reporting on collected data. Network managers and analysts can go to this report to select a system from a table, and assess its status in the following performance areas: CPU utilization, number of processes, memory utilization, page faults, swap utilization, and filesystem utilization.

Alternatively, the Snapshot report shows the same information, however, you select systems you want to view from a pick list when you invoke the report. For more information about selecting a system for a Snapshot report, see "Selecting Deferred Targets or Elements for a Snapshot Report" on page 1-27.

The QuickView and Snapshot reports comprise the following charts:

- The **System Selection List** table shows each system and its performance statistics based on the data polled for the last 6 hours. The statistic shown for

each metric is the average of all the data collected for that statistic over the 6-hour period.

System Selection List
Select device to see near real time reports

Device	CPU Utilization	Num of Processes	Swap Capacity	Avg Page Faults
wasp	30.58	159.54	30.83	55739.33
leyte	30.58	160.25	30.42	55199.79
nimitz	30.54	160.54	30.63	55760.63
yorktown	20.96	159.58	30.79	31294.29
laffey	20.92	161.38	31.04	31422.38
hornet	6.88	161.79	13.50	12416.75
eisenhower	5.63	157.79	11.92	12488.17
enterprise	5.58	160.58	12.08	12399.63

Double-click on a system in the table to update the content of the following drill-down components with performance data for the system:

- The **Number of Processes** line graph shows the number of processes for the selected system for each sample taken during the report period.

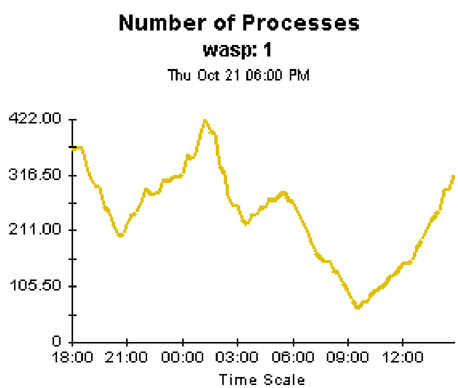

- The **CPU Utilization** line graph shows the selected system's CPU utilization for each sample taken during the report period.

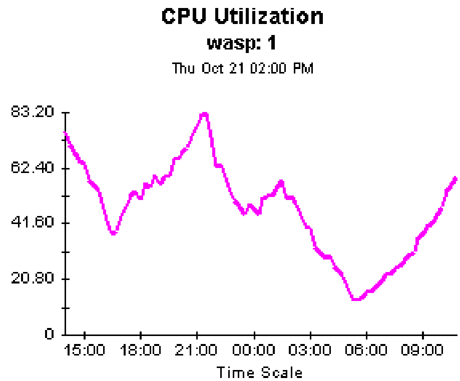

- The **Memory Utilization** line graph shows the selected system's memory utilization for each sample taken during the report period.

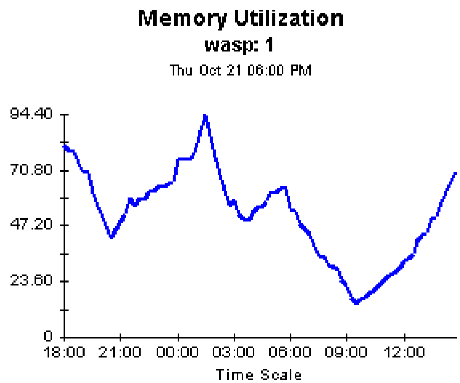

- The **Page Faults** line graph shows the selected system's page faults for each sample taken during the report period.

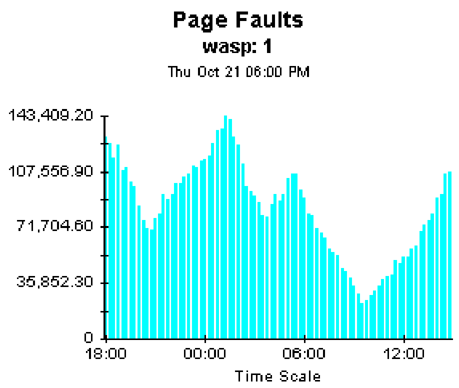

- The **Swap Utilization** line graph shows the selected system's swap utilization for each sample taken during the report period.

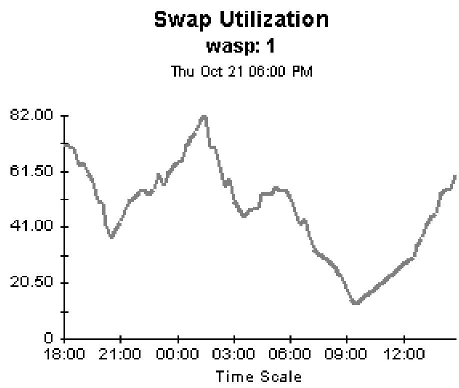

- The **Filesystem Utilization** line graph shows the selected system's filesystem utilization for each sample taken during the report period.

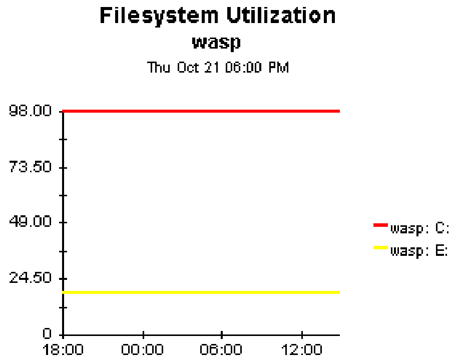

Capacity Planning--Filesystems Report

The Capacity Planning--Filesystems report shows the most overutilized and underutilized systems that are within 90 days of reaching a disk space utilization threshold. It allows the CIO and network manager to estimate utilization levels in 90 days. You can double-click on any interface in the Overutilized Systems or Underutilized Systems summary tables to display corresponding drill-down charts that detail daily data for Grade of Service; page faults; and CPU, memory, and swap utilization for the baseline period. It shows summaries of the filesystems that are forecasted to exceed 90% of the space used on the filesystem in the next 90 days for overutilized and less then 30% of the space used on the filesystem in the next 90 days for underutilized.

The Capacity Planning--Filesystems report comprises the following components:

- The **Overutilized System Filesystem** summary table shows the filesystems that are projected to be overutilized within the next 90-days. They are ranked in descending order by the 90-day projected filesystem utilization value. The filesystems using the most projected space within the next 90 days are shown.

All calculations are based on 95th percentile utilization values and are displayed as percentages in this table.

Overutilized Filesystems

Projected To Exceed 90% Utilization Within 90 Days

Device	Filesystem	Current Filesystem Utilization	Projected Util 30 Days	Projected Util 60 Days	Projected Util 90 Days
wasp	C:	98.00	98.00	98.00	98.00
yorktown	/export/home	95.00	95.00	95.00	95.00

Double-click on a system in this table to update the following components with performance data for the selection:

- The **Grade of Service** stacked bar chart shows GOS scores for CPU, memory, and swap utilization, and queue length over the baseline period for the selected system. See Table 12-7 for a description of the GOS weighting factors.

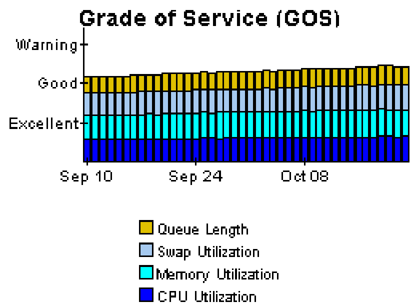

- The **Filesystem Utilization** line graph shows the average filesystem utilization over each day of the baseline period for the selected system.

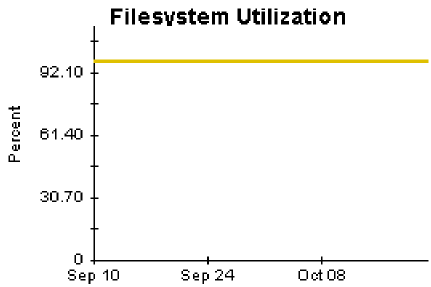

- The **CPU Utilization** line graph shows the average CPU utilization over each day of the baseline period for the selected system.

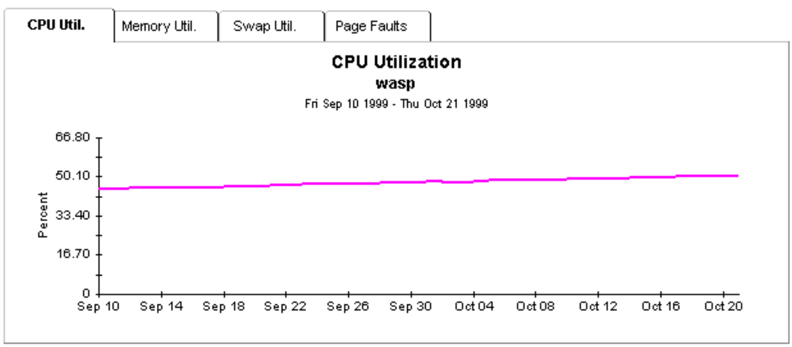

- The **Memory Utilization** line graph shows the average memory utilization over each day of the baseline period for the selected system.

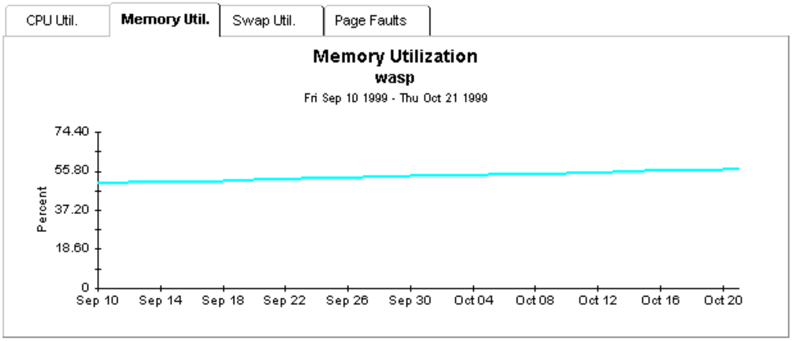

- The **Swap Utilization** line graph shows the average swap utilization over each day of the baseline period for the selected system.

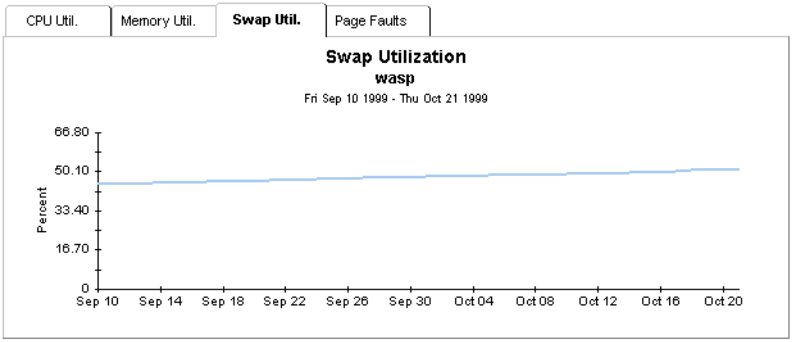

- The **Page Faults** bar chart shows the number of page faults for each day of the baseline period for the selected system.

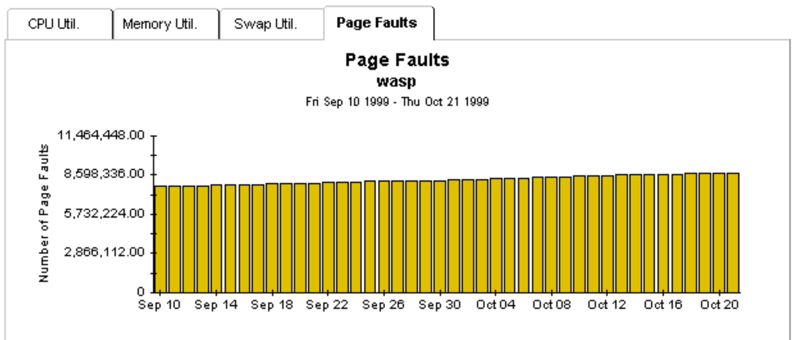

- The **Underutilized System Filesystem** summary table shows the systems that are projected to be the least utilized within the next 90 days. It is ranked in ascending sequence by the 90-day projected filesystem utilization value. All calculations are based on 95th percentile utilization values and are displayed as percentages in this table. Underutilized means that the system is furthest from reaching the filesystem utilization thresholds (30%) defined.

Underutilized Filesystems

Device	Filesystem	Current Filesystem Utilization	Projected Util 30 Days	Projected Util 60 Days	Projected Util 90 Days
yorktown	/export/home1	12.00	12.00	12.00	12.00
wasp	E:	18.00	18.00	18.00	18.00

Double-click on a system in this table to update the following components with performance data for the selection:

- The **Grade of Service** stacked bar chart shows GOS scores for queue length, and CPU, memory, and swap utilization over the baseline period for the selected system. See Table 12-7 for a description of GOS weighting factors.
- The **Filesystem Utilization** line graph shows the average filesystem utilization over each day of the baseline period for the selected system.

- The **CPU Utilization** line graph shows the average CPU utilization over each day of the baseline period for the selected system.
- The **Memory Utilization** line graph shows the average memory utilization over each day of the baseline period for the selected system.
- The **Swap Utilization** line graph shows the average swap utilization over each day of the baseline period for the selected system.
- The **Page Faults** bar chart shows the number of page faults for each day of the baseline period for the selected system.

Capacity Planning--Memory Report

The Capacity Planning--Memory report shows the most overutilized and underutilized systems that are within 90 days of reaching a memory utilization threshold. It allows the CIO and network manager to estimate utilization levels in 90 days. You can double-click on any interface in the Overutilized Systems or Underutilized Systems summary tables to display corresponding drill down charts that detail daily data for Grade of Service; page faults; and CPU, memory, and swap utilization for the baseline period. It shows summaries of the systems that are forecasted to exceed 100% of the memory used on the system in the next 90 days for overutilized and less then 30% of the memory used on the system in the next 90 days for underutilized.

Note: The charts for the Capacity Planning--Memory report are identical to the charts for the Capacity Planning--Filesystems report. For chart illustrations, see "Capacity Planning--Filesystems Report" on page 12-18.

The Capacity Planning--Memory report comprises the following components:

The **Overutilized System Memory** summary table shows the computer systems that are projected to be overutilized within the next 90-days ranked in descending order by the 90-day projected memory utilization value. All calculations are based on 95th percentile utilization values and are displayed as percentages in this table. Table 12-6 describes how memory utilization is computed.

Double-click on a system in this table to update the following components with performance data for the selection:

- The **Grade of Service** stacked bar chart shows GOS scores for CPU, memory, and swap utilization, and queue length over the baseline period for the selected system. See Table 12-7 for a description of the GOS weighting factors.
- The **Memory Utilization** line graph shows the average memory utilization over each day of the baseline period for the selected system.
- The **CPU Utilization** line graph shows the average CPU utilization over each day of the baseline period for the selected system.
- The **Filesystem Utilization** line graph shows the average filesystem utilization over each day of the baseline period for the selected system.
- The **Swap Utilization** line graph shows the average swap utilization over each day of the baseline period for the selected system.
- The **Page Faults** bar chart shows the number of page faults for each day of the baseline period for the selected system.

- The **Underutilized System Memory** summary table shows the systems that are projected to be the least utilized within the next 90-days ranked in ascending sequence by the 90-day projected memory utilization value. All calculations are based on 95th percentile utilization values and are displayed as percentages in this table. Underutilized means that the system is furthest from reaching the utilization thresholds (30%) defined.

  Double-click on a system in this table to update the following components with performance data for the selection:

  - The **Grade of Service** stacked bar chart shows GOS scores for queue length, and CPU, memory, and swap utilization over the baseline period for the selected system. See Table 12-7 for a description of GOS weighting factors.
  - The **Memory Utilization** line graph shows the average memory utilization over each day of the baseline period for the selected system.
  - The **CPU Utilization** line graph shows the average CPU utilization over each day of the baseline period for the selected system.

- The **Filesystem Utilization** line graph shows the average filesystem utilization over each day of the baseline period for the selected system.
- The **Swap Utilization** line graph shows the average swap utilization over each day of the baseline period for the selected system.
- The **Page Faults** bar chart shows the number of page faults for each day of the baseline period for the selected system.

Hot Spots Report

The Hot Spots report identifies specific areas of concern where threshold values have been exceeded. The top ten offenders, ranked by total number of exceptions, are presented in a summary table to assist the network management team in analyzing the problems. Each exception is automatically hot-linked to multiple drill down charts including grade of service, CPU utilization, memory utilization, swap capacity, and number of swaps. An exception detail table reveals the times and offending levels for each exception.

The Hot Spots report comprises the following components:

- The **Problem Summary for the Day** summary table displays the top ten systems in descending order by total number of exceptions (CPU + memory + total number of swaps + paging +page fault exceptions) for the report day. An exception occurs when an interface exceeds its defined CPU utilization (50%), memory utilization (80%), total number of swaps (>10), total number of pages

(>30,000), or total number of page faults (>100,000) threshold in any sample taken during the report day.

Problem Summary for the Day
System With Most Exceptions For The Day
Thu Oct 21 1999

Device	Interface	CPU Util Exceptions	Memory Util Exceptions	Paging Exceptions	Page Fault Exceptions	Swap Exceptions	Total Exceptions
leyte	1	54.00	18.00	75.00	38.00	13.00	198.00
wasp	1	54.00	15.00	75.00	40.00	12.00	196.00
nimitz	1	55.00	17.00	74.00	38.00	10.00	194.00
laffey	1	14.00	22.00	67.00	0.00	49.00	152.00
yorktown	1	15.00	18.00	68.00	0.00	51.00	152.00

Double-click on any system listed in the table to update the following components with performance data for the selection:

- The **Grade of Service** stacked bar chart shows weighted values for memory, swap, and CPU utilization, and queue length for the selected system over each hour on the report day. Table 12-7 gives the GOS weighting factors.

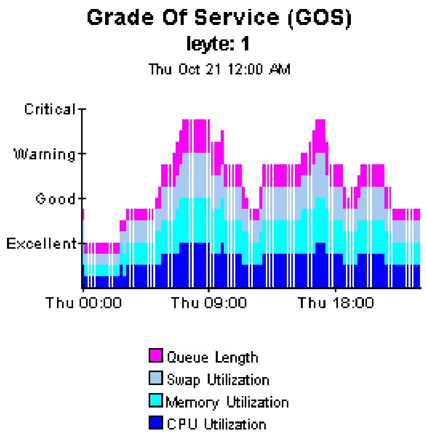

- The **CPU Utilization** line graph shows the percentage of CPU utilization for each hour of the report day for the specified system.

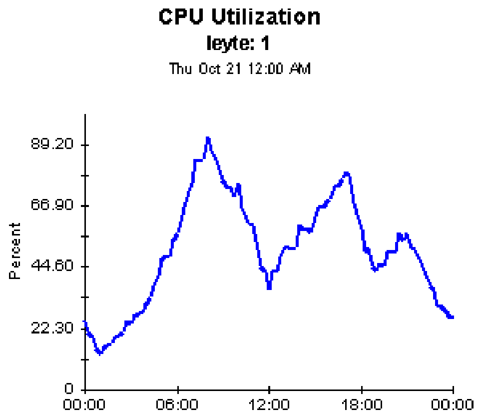

- The **Memory Utilization** line graph shows the average percentage of memory utilization for each hour of the report day for the specified system.

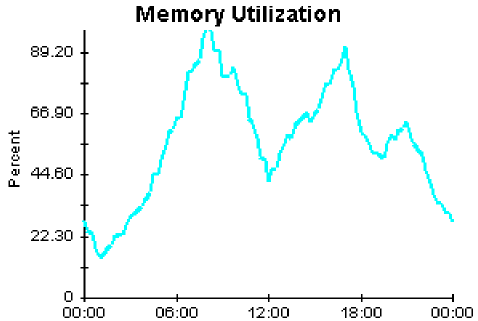

- The **Page Faults** bar chart shows the number of page faults for each hour of the report day for the specified system.

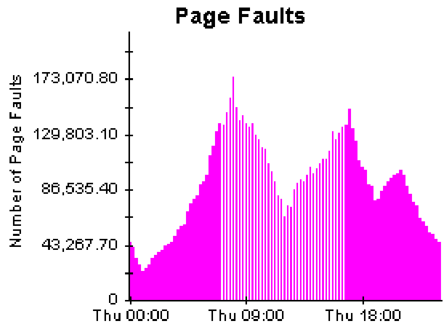

- The **Swap Capacity** line graph shows the average percentage of swap capacity available for each hour of the report day for the specified system.

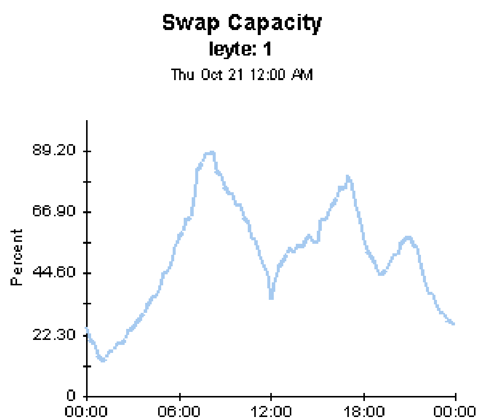

- The **Total Swaps** bar chart shows the total swaps for each hour of the report day for the specified system.

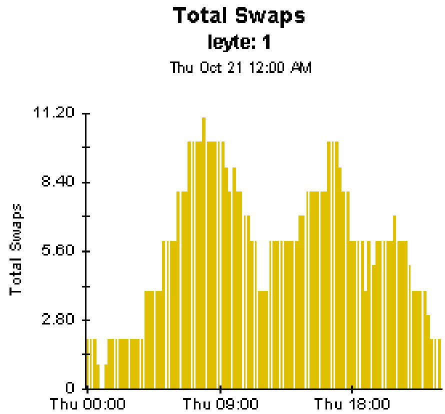

- The **Exception Details** table shows the time the sample was taken for each CPU and memory utilization, page fault, paging, and swapping exceptions and the values for the offending conditions. Table 12-6 describes how each metric is calculated.

Exception Details
leyte: 1
Thu Oct 21 04:15 AM

Time Period	Device	Interface	Description	CPU Utilization	Memory Utilization	Page Faults	Paging	Swa
Thu Oct 21 04:15 AM	leyte	1		34.00	38.50	59480.00	31174.00	4.
Thu Oct 21 04:30 AM	leyte	1		38.00	44.57	69721.00	34468.00	4.
Thu Oct 21 04:45 AM	leyte	1		41.00	44.81	74868.00	36068.00	4.
Thu Oct 21 05:00 AM	leyte	1		47.00	49.88	79117.00	38724.00	4.
Thu Oct 21 05:15 AM	leyte	1		49.00	53.23	81705.00	41937.00	6.
Thu Oct 21 05:30 AM	leyte	1		48.00	59.88	90874.00	44029.00	6.
Thu Oct 21 05:45 AM	leyte	1		54.00	60.93	93370.00	48035.00	6.
Thu Oct 21 06:00 AM	leyte	1		57.00	65.46	98659.00	49329.00	6.
Thu Oct 21 06:15 AM	leyte	1		61.00	65.23	113519.00	55762.00	8.

Top Ten Report

The Top Ten report lets the network manager know at-a-glance which systems produce the greatest throughput and exhibit the poorest health. The Top Ten report identifies the top ten systems with highest ranked metric values. The top ten elements, ranked by total volume and grade of service are listed by highest rank as well as highest change of rank. In addition to the ranks of the elements, the volume and utilization values are displayed for reference.

The Top Ten report comprises the following components:

- The **Systems with Highest Volume** summary table ranks the selected systems in ascending sequence by volume rank for the day.

Systems With Highest Volume
Thu Oct 21 1999

Device	Element	Volume	Volume Rank	Previous R
yorktown	2	12195099049.85	1	1
eisenhower	1	4567909937.96	2	3
nimitz	1	4560733456.28	3	2
yorktown	1	1225213583.37	4	5
laffey	1	1221760234.58	5	4
nimitz	2	1144531469.83	6	6
leyte	1	1137775458.31	7	7
eisenhower	2	459146832.48	8	10
hornet	1	457219636.28	9	8
wasp	1	457197063.48	10	9

- The **Systems with Greatest Change in Volume** summary table ranks the selected systems in ascending sequence by change in volume rank from the

previous day. A large change in rank identifies a volatile resource, and you should investigate the reasons for the high change in volume.

Systems With Greatest Change In Volume
Thu Oct 21 1999

Device	Element	Volume	Previous Volume	Volume Change Rank	Volume Rank	Previous Rank
yorktown	2	12195099049.85	12288955260.65	1	1	1
eisenhower	1	4567909937.96	4553401621.26	2	2	3
eisenhower	2	459146832.48	453566131.55	3	8	10
leyte	1	1137775458.31	1141841031.46	4	7	7
laffey	1	1221760234.58	1225594036.38	5	5	4
hornet	1	457219636.28	459759653.89	6	9	8
enterprise	1	407859790.15	409724021.40	7	11	11
yorktown	1	1225213583.37	1223440369.62	8	4	5
wasp	1	457197063.48	457569041.87	9	10	9
nimitz	2	1144531469.83	1144241107.29	10	6	6

- The **Systems with Highest Grade of Service Scores** summary table ranks the selected systems in ascending sequence by grade of service rank. It shows the grade of service score, grade of service rank, previous rank, daily averages for queue length, and CPU, memory, and swap utilization for the top ten device and element pairs.

Systems With Highest Grade of Service (GOS) Scores
Thu Oct 21 1999

Device	Element	GOS Score	GOS Rank	Previous Rank	CPU Utilization	Memory Utilization	Swap Utilization	Queue Lengt
nimitz	1	2.47	1	1	51.11	56.56	50.77	1.84
wasp	1	2.45	2	2	51.01	56.47	50.78	1.83
leyte	1	2.45	3	3	51.02	56.79	50.69	1.77
yorktown	1	2.25	4	4	34.21	58.63	50.48	1.54
laffey	1	2.24	5	5	34.09	58.51	50.47	1.59
hornet	1	1.27	6	6	11.66	27.69	22.32	0.69
eisenhower	1	1.18	7	7	9.40	23.46	19.85	0.00
enterprise	1	1.17	8	8	9.46	23.28	19.90	0.00

- The **Systems with the Greatest Change in Grade of Service Scores** summary table ranks the selected systems in ascending sequence by change in grade of service scores from the previous day. A large change in rank identifies

a volatile resource, and you should investigate the reasons for the high change in volume.

Systems With Greatest Change In Grade of Service (GOS) Scores
Thu Oct 21 1999

Device	Element	GOS Score	Previous GOS Score	GOS Change Rank	GOS Rank	Previous Rank	CPU Utilization	Me Utili
leyte	1	2.45	2.43	1	3	3	51.02	5
yorktown	1	2.25	2.26	2	4	4	34.21	5
hornet	1	1.27	1.26	3	6	6	11.66	2
nimitz	1	2.47	2.46	4	1	1	51.11	5
laffey	1	2.24	2.24	5	5	5	34.09	5
eisenhower	1	1.18	1.18	6	7	7	9.40	2
enterprise	1	1.17	1.18	7	8	8	9.46	2
wasp	1	2.45	2.45	8	2	2	51.01	5

Service Level Management Report

The Service Level Management report lets CIOs, CFOs, network managers, and customers know if their systems are meeting contracted service levels for availability and response times by reporting system availability across the enterprise over time and for individual Systems for the day. It reports the ten systems with lowest availability for the day and the ten systems with the highest network response times. The System Service Level Management report provides drilldowns for daily availability and daily and hourly response times for these elements.

Note: For an in-depth analysis of how to interpret the System Service Level Management report and for chart illustrations, see "Service Level Management Report" on page 3-41.

The Service Level Management report comprises the following components:

- The **Availability** list shows the systems with the lowest availability for the report day.

Double-click on a system in this list to update the following component with performance data for the selected system:

- The **Daily Availability** line graph shows the systems with the lowest availability for the report day compared to the service level for the baseline period.
- The **Hourly Availability** line graph shows the systems with the lowest availability for each hour of the report day.

- The **Response Time** list shows the systems with the highest response time for the report day.

  Double-click on a system in this list to update the following component with performance data for the selected system:

  - The **Daily Response Time** line graph shows the daily response time for the selected System for each sample over the report period (42 days).
  - The **Hourly Response Time** line graph shows the response time for the selected system for each sample for the preceding day.

Data Source

Source data for the System ReportPack is collected from the Empire Technology's SystemEDGE product. The data comes from four tables in the Empire MIB and one MIBlet from the MIB-II MIB. See the *TREND User's Guide* for more information about MIBlets.

Table 12-1 lists the variables from the kernelperf table in the empire.mib file.

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	OID String
cpu1Min	1.3.6.1.4.1.546.1.1.7.8.1
The percentage of time over the last minute the CPU was not idle. The agent continuously samples the CPU and updates this variable. On multi-processor machines, this object represents the value for the first CPU only. The Host Resources MIB hrProcessorTable breaks out the CPU load on a per-processor basis. This value is meaningless until at least two samples have been taken. On NT systems, this value is the average over all CPUs.	
cpu5Min	1.3.6.1.4.1.546.1.1.7.8.2
The percentage of time over the last 5 minutes the CPU was not idle. The agent continuously samples the CPU and updates this variable. On multi-processor machines, this object represents the value for the first CPU only. The Host Resources MIB hrProcessorTable breaks out the CPU load on a per-processor basis. This value is meaningless until at least two samples have been taken. On NT systems, this value is the average over all CPUs.	
cpu15Min	1.3.6.1.4.1.546.1.1.7.8.3
The percentage of time over the last 15 minutes the CPU was not idle. The agent continuously samples the CPU and updates this variable. On multi-processor machines, this object represents the value for the first CPU only. The Host Resources MIB hrProcessorTable breaks out the CPU load on a per-processor basis. This value is meaningless until at least two samples have been taken. On NT systems, this value is the average over all CPUs.	
runQlen	1.3.6.1.4.1.546.1.1.7.8.4

(1 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	OID String
The length of the scheduler's run queue sampled over a fixed period of time. The sampling interval is system dependent. In SunOS, the kernel updates this parameter once every 5 seconds. The system's load average is computed from the value of runQLen; the load average is defined as the average runQLen value over some interval.	
diskWaitNum	1.3.6.1.4.1.546.1.1.7.8.5
The number of jobs that are waiting on disk I/O sampled over a fixed period of time. The sampling interval is system dependent. In SunOS, the kernel updates this parameter once every 5 seconds. A large value may indicate an overloaded system although larger values may be observed on file server systems.	
pageWaitNum	1.3.6.1.4.1.546.1.1.7.8.6
The number of jobs that are waiting on page I/O sampled over a fixed period of time. The sampling interval is system dependent. In SunOS, the kernel updates this parameter once every 5 seconds. A large value may indicate the system does not have enough memory for the current load.	
swapActive	1.3.6.1.4.1.546.1.1.7.8.7
The number of jobs that are active yet swapped out of memory, sampled over a fixed period of time. The sampling interval is system dependent. In SunOS, the kernel updates this parameter once every 5 seconds. A high value for this object may indicate that system memory is not sufficient for the current load.	
sleepActive	1.3.6.1.4.1.546.1.1.7.8.8

(2 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	**OID String**
The number of jobs that are active yet sleeping in core, sampled over a fixed period of time. The sampling interval is system dependent. In SunOS, the kernel updates this parameter once every 5 seconds. A process is active yet sleeping in core when it is not running and not swapped out to secondary storage (swap).	
memInUse	1.3.6.1.4.1.546.1.1.7.8.9
The total real KBytes of memory in use sampled over a fixed period of time. The sampling interval is system dependent. In SunOS, the kernel updates this parameter once every 5 seconds. Comparing memInUse to memory (in the Empire ’system’ group) can help indicate if a system has sufficient memory for its current load.	
activeMem	1.3.6.1.4.1.546.1.1.7.8.10
The total active real memory (KBytes) sampled over a fixed period of time. The sampling interval is system dependent. In SunOS, the kernel updates this parameter once every 5 seconds. Memory is active when it is actually being used by a process running in memory. This value should be less than or equal to memInUse. A high value may indicate a system with insufficient memory.	
numProcs	1.3.6.1.4.1.546.1.1.7.8.11
The number of processes table slots currently allocated and in use. This number equals the number of currently running processes on the system. A high value may indicate an overloaded system although higher values should be expected on servers and systems shared by many users.	

(3 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	OID String
numOpenFiles	1.3.6.1.4.1.546.1.1.7.8.12
The system-wide total number of open files.	
swapInUse	1.3.6.1.4.1.546.1.1.7.8.13
The amount of swap space (KBytes) currently in use by the system. Swap space is used when memory is in short supply. Typically, entire processes are swapped out to secondary storage. Comparing this value to totalSwap (in Empire's kernelConfig group) can help indicate if additional swap space is needed.	
numSwitches	1.3.6.1.4.1.546.1.1.7.8.14
The total number of context switches that have occurred since the kernel was last initialized. A context switch occurs each time a process gives up the CPU and another takes its place. This counter reflects the level of system activity; a high rate of context switching is indicative of the system load.	
numTraps	1.3.6.1.4.1.546.1.1.7.8.15
The total number of traps switching a process from user to kernel mode that have occurred since the kernel was last initialized. Traps are hardware exceptions that occur and are usually caused by the currently running process. Example traps include those for incorrect arithmetic operations, segmentation violations or page faults. An excessive trap rate can indicate a faulty program, excessive system loading, or some other abnormality.	
numSyscalls	1.3.6.1.4.1.546.1.1.7.8.16

(4 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	**OID String**
The total number of system calls that have occurred since the kernel was last initialized. System calls occur when a process calls a function that must execute in kernel or privileged mode; it is a indicator of the system load and may be higher on servers or machines shared by many users.	
numInterrupts	1.3.6.1.4.1.546.1.1.7.8.17
The total number of device interrupts that have occurred since the kernel was last initialized. Devices and peripherals interrupt the operating system when they require the attention of the system. An excessive rate of interrupts can indicate a heavily loaded machine; servers and machines shared by many users will typically have a higher rate though.	
numPageSwapIns	1.3.6.1.4.1.546.1.1.7.8.18
The total number of pages that have been swapped in since the kernel was last initialized. Page swapping occurs when a process is swapped to and from secondary storage. A high rate of page swap activity can indicate an overloaded system.	
numPageSwapOuts	1.3.6.1.4.1.546.1.1.7.8.19
The total number of pages that have been swapped out since the kernel was last initialized. Page swapping occurs when a process is swapped to and from secondary storage. A high rate of page swap activity can indicate an overloaded system.	
numSwapIns	1.3.6.1.4.1.546.1.1.7.8.20

(5 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	**OID String**
The total number of processes that have been swapped in since the kernel was last initialized. Processes are swapped to and from secondary storage when memory is in short supply. A high rate may indicate an overloaded system.	
numSwapOuts	1.3.6.1.4.1.546.1.1.7.8.21
The total number of processes that have been swapped out since the kernel was last initialized. Processes are swapped to and from secondary storage when memory is in short supply. A high rate may indicate an overloaded system.	
numPageIns	1.3.6.1.4.1.546.1.1.7.8.22
The total number of pages that have been paged in since the kernel was last initialized. Page-in is an operation performed by the virtual memory system in which the contents of a page are read from secondary storage. A high rate of change may indicate an overloaded system.	
numPageOuts	1.3.6.1.4.1.546.1.1.7.8.23
The total number of pages that have been paged out since the kernel was last initialized. Page-out is an operation performed by the virtual memory system in which the contents of a page are written to secondary storage. A high rate of change may indicate an overloaded system.	
numPageReclaims	1.3.6.1.4.1.546.1.1.7.8.24

(6 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	**OID String**
The total pages that have been reclaimed from the freelist since the kernel was last initialized. Memory pages are stored on a free-list when not in use. If a page has been acquired from the free-list, it has been reclaimed. A high rate of changed indicates an active system which is not generally considered to be overloaded.	
numPageFaults	1.3.6.1.4.1.546.1.1.7.8.25
The total number of page faults that have occurred since the kernel was last initialized. A page fault occurs when a running process attempts to access a virtual memory page that is not currently in physical memory resulting in a hardware page fault. A high rage of page faults may indicate an overloaded system or one with insufficient memory.	
loadAverage1Min	1.3.6.1.4.1.546.1.1.7.8.26
The load average over the last 1 minute. The load average represents the average number of jobs in the run queue over the specified time range. The value reported is the load average multiplied by 100. For example, the value 50 represents a load average of 0.50.	
loadAverage5Min	1.3.6.1.4.1.546.1.1.7.8.27
The load average over the last 5 minutes. The load average represents the average number of jobs in the run queue over the specified time range. The value reported is the load average multiplied by 100. For example, the value 50 represents a load average of 0.50.	
loadAverage15Min	1.3.6.1.4.1.546.1.1.7.8.28

(7 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	**OID String**
The load average over the last 15 minutes. The load average represents the average number of jobs in the run queue over the specified time range. The value reported is the load average multiplied by 100. For example, the value 50 represents a load average of 0.50.	
totalSwapSpace	1.3.6.1.4.1.546.1.1.7.8.29
Total system swap space in KBytes. This object is identical to the kernelConfig.totalSwap(12) object and is repeated here for the convenience of management stations which cannot simultaneously poll groups of objects in different branches of the MIB.	
swapCapacity	1.3.6.1.4.1.546.1.1.7.8.30
The percentage of the system's total swap in use. The value of this object can be computed by dividing swapInUse(13) by totalSwapSpace(29) and multiplying by 100. It is included here for convenient polling, monitoring, and history sampling. This object is an aggregate value over all the swap partitions and areas.	
memCapacity	1.3.6.1.4.1.546.1.1.7.8.31
The percentage of the system's memory in use. The value of this object can be approximated by dividing activeMem (10) by memory (system.3) and converting to a percentage. Because virtual memory allows systems to operate with a higher memCapacity value, this value alone should not be used to judge the amount of the system load. Managers should also monitor swapCapacity.	
memInUseCapacity	1.3.6.1.4.1.546.1.1.7.8.32

(8 of 9)

Table 12-1: System Variables Polled from Empire kernelperf Table

Variable	OID String
The percentage of the system's memory in use. The value of this object can be approximated by dividing memInUse (9) by memory (system.3) and converting to a percentage. Because virtual memory allows systems to operate with a higher memCapacity value, this value alone should not be used to judge the amount of the system load. Managers should also monitor swapCapacity.	

(9 of 9)

The MIB II server_volume table is a miblet of the MIB II ifEntry table. Table 12-2 lists the variables.

Table 12-2: System Variables Polled from MIB II Server Volume Table

Variable	OID String
ifSpeed	1.3.6.1.2.1.2.2.1.5
An estimate of the interface's current bandwidth in bits per second. For interfaces that do not vary in bandwidth or for those where no accurate estimation can be made, this object should contain the nominal bandwidth.	
ifInOctets	1.3.6.1.2.1.2.2.1.10
The total number of octets received on the interface, including framing characters.	
ifOutOctets	1.3.6.1.2.1.2.2.1.16

(1 of 2)

Table 12-2: System Variables Polled from MIB II Server Volume Table

Variable	OID String
The total number of octets transmitted out of the interface, including framing characters.	

(2 of 2)

Table 12-3 lists the variables from the cpuStatsEntry table in the empire.mib file.

Table 12-3: System Variables Polled from Empire cpuStatsEntry Table

Variable	OID String
cpuStatsIndex	1.3.6.1.4.1.546.13.1.1.1
An index that uniquely identifies an entry in the cpuStatsTable.	
cpuStatsDescr	1.3.6.1.4.1.546.13.1.1.2
The type of CPU for which statistics are reported. On Solaris 2.x, this description also includes the CPU’s clock speed.	
cpuStatsIdle	1.3.6.1.4.1.546.13.1.1.3
The total number of ’ticks’ spent by this CPU in Idle mode. This variable is not supported by NT. This value may wrap a 32-bit counter.	
cpuStatsUser	1.3.6.1.4.1.546.13.1.1.4
The total number of ’ticks’ spent by this CPU in User mode. This value may wrap a 32-bit counter.	

(1 of 3)

Table 12-3: System Variables Polled from Empire cpuStatsEntry Table

Variable	**OID String**
cpuStatsSys	1.3.6.1.4.1.546.13.1.1.5
The total number of ’ticks’ spent by this CPU in Kernel or system mode. This value may wrap a 32-bit counter.	
cpuStatsWait	1.3.6.1.4.1.546.13.1.1.6
The total number of ’ticks’ spent by this CPU in Wait mode. Wait mode is that time spent waiting on I/O or swap. It may wrap a 32-bit counter.	
cpuStatsLastUpdate	1.3.6.1.4.1.546.13.1.1.7
The time (based on sysUpTime) that the counters for this CPU were last updated. Because these values are cached periodically by the agent, the time of last update is important for determining the interval over which CPU mode percentages are calculated.	
cpuStatsIdlePercent	1.3.6.1.4.1.546.13.1.1.8
The percentage of time (over the sample period) the system’s CPU was idle.	
cpuStatsUserPercent	1.3.6.1.4.1.546.13.1.1.9
The percentage of time (over the sample period) the system’s CPU was executing in user mode.	
cpuStatsSysPercent	1.3.6.1.4.1.546.13.1.1.10

(2 of 3)

Table 12-3: System Variables Polled from Empire cpuStatsEntry Table

Variable	OID String
The percentage of time (over the sample period) that the system's CPU was executing the kernel or operating system.	
cpuStatsWaitPercent	1.3.6.1.4.1.546.13.1.1.11
The percentage of time (over the sample period) the system's CPU was waiting for I/O. This object may not be supported by the underlying operating system.	

(3 of 3)

Table 12-4 lists the variables from the diskStatsEntry table in the empire.mib file.

Table 12-4: System Variables Polled from Empire diskStatsEntry Table

Variable	OID String
diskStatsIndex	1.3.6.1.4.1.546.12.1.1.1
An index that uniquely identifies an entry in the diskStatsTable.	
diskStatsQueueLength	1.3.6.1.4.1.546.12.1.1.2
The average number of operations waiting in the disk's service queue over the last measurement period.	
diskStatsServiceTime	1.3.6.1.4.1.546.12.1.1.3

(1 of 3)

Table 12-4: System Variables Polled from Empire diskStatsEntry Table

Variable	**OID String**
The average service time in milliseconds for operations served on this disk over the last measurement period. This could be expressed as: disk-busy-time / number-of-transfers.	
diskStatsUtilization	1.3.6.1.4.1.546.12.1.1.4
The utilization rate (percentage utilization) for this disk over the last measurement period. This could also be expressed as: (disk-busy-time / elapsed-time) * 100.	
diskStatsKBytesTransferred	1.3.6.1.4.1.546.12.1.1.5
The number of KBytes transferred to/from this disk.	
diskStatstransfers	1.3.6.1.4.1.546.12.1.1.6
The number of transfer operations to/from this disk.	
disksStatsReads	1.3.6.1.4.1.546.12.1.1.7
The number of read operations from this disk.	
diskStatsWrites	1.3.6.1.4.1.546.12.1.1.8
The number of write operations to this disk.	
diskStatsHostmibDevTableIndexx	1.3.6.1.4.1.546.12.1.1.9

(2 of 3)

Table 12-4: System Variables Polled from Empire diskStatsEntry Table

Variable	OID String
The index of this disk device in the device table of the Host Resources MIB.	
diskStatsLastUpdate	1.3.6.1.4.1.546.12.1.1.10
The time in SNMP Time Ticks of the last update of this diskStatsTable entry.	

(3 of 3)

Table 12-5 lists the variables from the devTableEntry table in the empire.mib file.

Table 12-5: System Variables Polled from Empire devTableEntry Table

Variable	OID String
devIndex	1.3.6.1.4.1.546.1.1.1.7.1.1
This object indicates the row instance. The instance is derived from the index number from the mounted file systems file.	
devDevice	1.3.6.1.4.1.546.1.1.1.7.1.2
The device name (e.g. /dev/sd0a) corresponding to the mounted file system. For example, the '/' file system is commonly mounted from the device '/dev/sd0a'. For network file systems, the value corresponds to the remote host and remote file system name (e.g. emptech:/usr/local).	
devMntPt	1.3.6.1.4.1.546.1.1.1.7.1.3

(1 of 4)

Table 12-5: System Variables Polled from Empire devTableEntry Table

Variable	**OID String**
The mount point or path name in the local file system where the device is mounted (e.g. /home).	
devBsize	1.3.6.1.4.1.546.1.1.1.7.1.4
The block size of the mounted device in bytes. File systems and disk devices generally allocate data in fixed size 'blocks' generally referred to as the block size of the device. Common block sizes are 1024 (1K) and 2048 (2K) bytes.	
devTblks	1.3.6.1.4.1.546.1.1.1.7.1.5
The total number of blocks contained in the mounted device.	
devFblks	1.3.6.1.4.1.546.1.1.1.7.1.6
The total number of free or unallocated blocks in the mounted device.	
devTfiles	1.3.6.1.4.1.546.1.1.1.7.1.7
Each disk device and file system allocates a fixed, maximum number of files that can be written to the device. The value of this object indicates the maximum number of files that can be written to this device.	
devFfiles	1.3.6.1.4.1.546.1.1.1.7.1.8
The number of free or unallocated files that can be created on this device.	
devMaxNameLen	1.3.6.1.4.1.546.1.1.1.7.1.9

(2 of 4)

Table 12-5: System Variables Polled from Empire devTableEntry Table

Variable	**OID String**
The maximum file name length that this mounted device supports. If no fixed limit exists, 0 is returned.	
devType	1.3.6.1.4.1.546.1.1.1.7.1.10
The type of the mounted device. If this underlying device does not support feature, 0 is returned.	
devFsid	1.3.6.1.4.1.546.1.1.1.7.1.11
The (possibly unique) file system ID of the mounted device. V.4 systems support its uniqueness while SunOS does not. For devices or operating systems not supporting this object, the object instance identifier is returned instead.	
devFstr	1.3.6.1.4.1.546.1.1.1.7.1.12
Some devices support a description string. If supported by the underlying device and operating system, that device specific description string is returned. If not supported, the string 'Not supported' is returned. An example description string may be 'User home directories'.	
devUnmount	1.3.6.1.4.1.546.1.1.1.7.1.13
Writing a delete(1) to this object causes the agent to unmount this device and remove its entry from the mounted file system table. Writing any other value to this object will cause no action to be taken by the agent.	
devCapacity	1.3.6.1.4.1.546.1.1.1.7.1.14

(3 of 4)

Table 12-5: System Variables Polled from Empire devTableEntry Table

Variable	OID String
The percentage of the device's total capacity in use.	

(4 of 4)

Basic Metric Calculations

The performance metrics used in the System reports are based on the calculations listed in Table 12-6.

Table 12-6: System Metrics

Metric	Computed As ...
Memory Utilization	((activemem/meminuse)*100)
Total Swaps	numswapins+numswapouts
Volume	inoctets + outoctets

Grade of Service Calculation

Table 12-7 lists how the metrics are weighted in the System Grade of Service (GOS) charts.

Table 12-7: System GOS Calculations

GOS Calculation	CPU Utilization	Memory Utilization	Swap Utilization	Queue Length
GOS weight factor	25%	25%	25%	25%
GOS score 1 (Excellent)	0-24%	0-24%	0-24%	0-1
GOS score 2 (Good)	25-49%	25-49%	25-49%	2 < 3
GOS score 4 (Warning)	50-74%	50-74%	50-74%	3-4
GOS score 5 (Critical)	75-100%	75-100%	75-100%	5 or more

For more information about the Grade of Service metric, see Table 2-2 on page 2-6. For more information about Grade of Service as it applies to each type of report, see “Analyzing TREND Reports” on page 3-1.

TREND

TREND

A Property File Formats

This appendix describes the property file formats for the following ReportPacks:

- LAN/WAN Connectivity
- Frame Relay
- Router
- Cisco Router
- Bay Router

Note: The Cisco and Bay Router ReportPacks use the same property file formats as the Router ReportPack. Ensure that you have the Router ReportPack installed before you use the Cisco or Bay Router ReportPacks.

You generate or create property files when you use the property export feature. An example property file is provided for each ReportPack. It is located in the following directory:

```
{DPIPE_HOME}/packages/<ReportPack>/<ReportPack>.ap/
PropertyData/Archive
```

When you create your own property files, ensure you adhere to the formats specified for each ReportPack. For more information about the property import feature, see "Using the Property Import Feature" on page 1-14.

Note: Do not change the dsi_target_name or dsi_table_key columns if you used the property export feature to create the data files. If you are creating this file manually, you will need to provide them.

LAN/WAN Connectivity ReportPack

Table A-1 lists the property file formats for the LAN/WAN Connectivity ReportPack.

Table A-1: LAN/WAN Connectivity ReportPack Property File Formats

Column Name	Function	Value Type
`dsi_target_name`	Device name, which is the name you use to poll the device.	varchar (64)
`dsi_table_key`	The device index, which is the ifindex value for the interface.	varchar (128)
`IfSpeed`	The interface speed in bits/second. The ReportPack requires this value.	float
`cust_id`	Unique integer value for each customer.	integer

Table A-1: LAN/WAN Connectivity ReportPack Property File Formats

Column Name	Function	Value Type
`customer_name`	Customer name associated with cust_id.	varchar (128)
`region_id`	Unique integer value for each region.	integer
`region_name`	Region name associated with region_id value.	varchar (128)
`dsi_descr`	A description of the interface.	varchar (128)

Frame Relay ReportPack

Table A-2 lists the property file formats for the Frame Relay ReportPack.

Table A-2: Frame Relay ReportPack Property File Formats

Column Name	Function	Value Type
`dsi_target_name`	Device name, which is the name you use to poll the device.	varchar (64)
`dsi_table_key`	The device index, which is the ifindex value of the interface and the DLCI value of the PVC concatenated and separated by a period, for example, 3.234.	varchar (128)
`port_speed`	The port speed in bits/second. The ReportPack requires this value.	float

Table A-2: Frame Relay ReportPack Property File Formats

Column Name	Function	Value Type
pvc_cir	The PVC's Committed Information Rate in bits/second. The ReportPack requires this value.	float
cust_id	Unique integer value for each customer.	integer
customer_name	Customer name associated with cust_id.	varchar (128)
region_id	Unique integer value for each region.	integer
region_name	Region name associated with region_id value.	varchar (128)
dsi_descr	A description of the PVC.	varchar (128)

Router ReportPack

Table A-3 lists the property file formats for the Router ReportPack.

Table A-3: Router ReportPack Property File Formats

Column Name	Function	Value Type
dsi_target_name	Device name, which is the name you use to poll the device.	varchar (64)
dsi_table_key	The device index, which is 0.	varchar (128)

Table A-3: Router ReportPack Property File Formats

Column Name	Function	Value Type
`cust_id`	Unique integer value for each customer.	integer
`customer_name`	Customer name associated with cust_id.	varchar (128)
`region_id`	Unique integer value for each region.	integer
`region_name`	Region name associated with region_id value.	varchar (128)
`dsi_descr`	A description of the router.	varchar (128)

TREND

TREND

Index

D

E

F

Index

N

P

Q

R

S

Index

T

U

V

W